

Data Stream Algorithms for Codeword Testing*

Atri Rudra Steve Uurtamo

Department of Computer Science and Engineering,
University at Buffalo, The State University of New York,
Buffalo, NY, 14620.
{atri,uurtamo}@buffalo.edu

Abstract

Motivated by applications in storage systems and property testing, we study data stream algorithms for local testing and tolerant testing of codes. Ideally, we would like to know whether there exist asymptotically good codes that can be local/tolerant tested with one-pass, poly-log space data stream algorithms.

We show that for the error detection problem (and hence, the local testing problem), there exists a one-pass, log-space data stream algorithm for a broad class of asymptotically good codes, including the Reed-Solomon (RS) code and expander codes. In our technically more involved result, we give a one-pass, $O(e \log^2 n)$ -space algorithm for RS (and related) codes with dimension k and block length n that can distinguish between the cases when the Hamming distance between the received word and the code is at most e and at least $a \cdot e$ for some absolute constant $a > 1$. For RS codes with random errors, we can obtain $e \leq O(n/k)$. For folded RS codes, we obtain similar results for worst-case errors as long as $e \leq (n/k)^{1-\varepsilon}$ for any constant $\varepsilon > 0$. These results follow by reducing the tolerant testing problem to the error detection problem using results from group testing and the list decodability of the code. We also show that using our techniques, the space requirement and the upper bound of $e \leq O(n/k)$ cannot be improved by more than logarithmic factors.

1 Introduction

In this work, we consider data stream algorithms for local testing and tolerant testing of error-correcting codes. The local testing problem for a code $C \subseteq \Sigma^n$ is the following: given a received word $\vec{y} \in \Sigma^n$, we need to figure out if $\vec{y} \in C$ or if \vec{y} differs from every codeword in C in at least $0 < e \leq n$ positions (i.e. the Hamming distance of \vec{y} from every $\vec{c} \in C$, denoted by $\Delta(\vec{y}, \vec{c})$, is at least e). If $e = 1$, then this is the error-detection problem. In the tolerant testing problem, given \vec{y} , we need to decide whether \vec{y} is at a distance at most e_1 from some codeword or if it has distance at least $e_2 > e_1$ from every codeword in C . Ideally, we would like to answer the following (see Section 2 for definitions related to codes):

Question 1. *Do there exist asymptotically good codes that can be (tolerant) tested by one-pass, poly-log space data stream algorithms?*

*Research supported by NSF CAREER Award CCF-0844796.

To the best of our knowledge, ours is the first work that considers this natural problem. We begin with the motivation for our work.

Property Testing. Local testing of codes has been extensively studied under the stricter requirements of property testing. Under the property testing requirements, one needs to solve the local testing problem by (ideally) only accessing a constant number of positions in \vec{y} . Codes that can be locally tested with a constant number of queries have been instrumental in the development of the PCP machinery, starting with the original proof of the PCP theorem [3, 4]. The current record is due to Dinur, who presents codes that have inverse poly-log rate, linear distance and can be locally tested with a constant number of queries [12].

General lower bounds on local testing of codes, however, have been scarce. (See, e.g. the recent paper [7]). In particular, it is not known if there are asymptotically good codes that can be locally tested with a constant number of queries. The question remains open even if one considers the harder task of tolerant testing with a constant number of *non-adaptive* queries [16].

It is not too hard to see that a non-adaptive tolerant tester that makes a constant number of queries gives a single pass, log-space data stream algorithm for tolerant testing. (See Section 3 for a proof.) Thus, if one could prove that any one-pass data stream algorithm for local/tolerant testing of asymptotically good codes requires $\omega(\log n)$ space, then they will have answered the question in the negative (at least for non-adaptive queries). This could present a new approach to attack the question of local/tolerant testing with a constant number of queries.

Next, we discuss the implications of *positive* results for local/tolerant testing.

Applications in Storage Systems. Codes are used in current day storage systems such as optical storage (CDs and DVDs), RAID ([11]) and ECC memory ([10]). Storage systems, up until recently, used simple codes such as the parity code and checksums, which have (trivial) data stream algorithms for error detection. However, the parity code cannot detect even two errors. With the explosion in the amount of data that needs to be stored, errors are becoming more frequent. This situation will become worse as more data gets stored on disks [14]. Thus, we need to use codes that can handle more errors.

Reed-Solomon (RS) codes, which are used widely in storage systems (e.g. in CDs and DVDs and more recently in RAID), are well-known to have good error correcting capabilities. However, the conventional error detection for RS codes is not space or pass efficient. Thus, a natural question to ask is if one can design a data stream algorithm to perform error detection for RS codes.

It would be remiss of us not to point out that unlike a typical application of a data stream algorithm where n is very large, in real life deployments of RS codes, n is relatively small. However, if one needs to implement the error detection algorithm in controllers on disks then it would be advantageous to use a data stream algorithm so that it is feasible to perform error detection with every read. Another way to use error detection is in *data scrubbing* [14]. In this scenario, during idle time or low activity periods, error detection is run on the entire disk to catch errors. In addition, the single pass requirement means that we will probe each bit on a disk only once, which is good for longevity of data. Finally, it would be helpful to complement an error-detection algorithm with a data stream algorithm that could also *locate* the errors.

It is also plausible that the efficiency of the data stream algorithms will make it feasible to use RS codes of block length (and alphabet size) considerably larger than the ones currently used in practice.

Before we delve into the description of our results, we would like to point out a few things.

First, for the storage application, designing algorithms for a widely used code such as the RS code will be more valuable than answering Question 1 in the affirmative via some new code. Second, it is known that for local testing of a RS code of dimension k , at least k queries need to be made in the property testing world¹. However, this does not rule out the possibility of a local/tolerant tester in the data stream world. Finally, for storage systems, solving the tolerant testing problem even for large constants e_1 and e_2 would be interesting.

Our Results. We give a one-pass, poly-log space, randomized algorithm to perform error detection for a broad class of asymptotically good codes such as Reed-Solomon (RS) and expander codes. As a complementary result, we also show that deterministic data stream algorithms (even with multiple passes) require linear space for such codes. Thus, for local testing we answer Question 1 in the affirmative. This should be contrasted with the situation in property testing, where it is known that for both asymptotically good RS and expander codes, a linear number of queries is required. (The lower bound for RS codes was discussed in the paragraph above and the result for expander codes follows from [8].)

It turns out that using existing results for tolerant testing of Reed-Muller codes over large alphabets [16], one can answer Question 1 in the affirmative for tolerant testing, though with $O(n^\varepsilon)$ space for any constant $\varepsilon > 0$. (See Section 3.2 for more details.)

Given the practical importance of RS codes and given the fact that local testing for RS codes with data stream constraints is possible, for the rest of the paper we focus mostly on tolerant testing of RS and related codes. We first remark that a naive tolerant testing algorithm for RS codes that can be implemented in $O(e \log n)$ space is to go through all the $\sum_{i=1}^e \binom{n}{i}$ possible error locations S and check if the received word projected outside of S belongs to the corresponding projected down RS code. (This works as long as $e \leq n - k$, which is true w.l.o.g. since $n - k$ is the covering radius of a RS code of dimension k and block length n .) Using our error detection algorithm for RS codes, this can be implemented as a one-pass $O(e \log n)$ -space data stream algorithm. Unfortunately, the running time of this algorithm is prohibitive, even for moderate values of e .

In this paper, we match the parameters above to within a log factor but with a (small) polynomial running time for values of e much larger than a constant. In particular, we present a one-pass, $O(e \log^2 n)$ -space, polynomial time randomized algorithm for a RS code C with dimension k and block length n that can distinguish between the cases when the Hamming distance between \vec{y} and C is at most e and at least $a \cdot e$ (for some constant $a > 1$). This reduction works when $e(e + k) \leq O(n)$. If we are dealing with random errors, then we can obtain $ek \leq O(n)$. Using known results on list decodability of folded RS codes [17], we obtain similar results for worst case errors for $e \leq (n/k)^{1-\varepsilon}$ for any constant $\varepsilon > 0$. As a byproduct, our algorithms also locate the errors (if the number of errors is bounded by e), which is desirable for a storage application. We also show that using our techniques, the space requirement and the upper bound of $e \leq O(n/k)$ cannot be improved by more than logarithmic factors.

Ideally, we would like our data stream algorithms to spend poly-log time per input position. However, in this paper we will tolerate polynomial time algorithms. In particular, naive implementations of the tolerant testing algorithms take $\tilde{O}(n^2)$ time. We also show that at the expense of slightly worse parameters, we can achieve a running time of $\tilde{O}(ne)$ for certain RS codes.

Our Techniques. It is well-known that error detection for any linear code can be done by checking if the product of the received word with the parity check matrix is the all zeros vector. We turn this

¹This follows from the fact that the “dual” code has distance $k + 1$.

into a one-pass low space data stream algorithm using the well-known finger printing method. The only difference from the usual fingerprinting method, where one uses any large enough field, is that we need to use a large enough extension field of the finite field over which the code is defined. To show the necessity of randomness, we use the well-known fooling set method from communication complexity [19]. However, unlike the usual application of two-party communication complexity in data stream algorithms, where the stream is broken up into two fixed portions, in our case we need to be careful about how we divide up the input. Details can be found in Section 5.

We now move on to our tolerant testing algorithm. We begin with the connection to group testing. Let \vec{c} be the closest codeword to \vec{y} and let $\vec{x} \in \{0,1\}^n$ denote the binary vector where $x_i = 1$ iff $y_i \neq c_i$. Now assume we could access \vec{x} in the following manner: pick a subset $Q \subseteq [n]$ and check if $\vec{x}_Q = \vec{0}$ or not (where \vec{x}_Q denotes \vec{x} projected onto indices in Q). Then can we come up with a clever way of non-adaptively choosing the tests such that at the end we know whether $\text{WT}(\vec{x}) \leq e$ or not? It turns out that we can use group testing to construct such an algorithm. In fact, using e -disjunct matrices (cf. [13]), we can design non-adaptive tests such that given the answers to the tests one could compute \vec{x} if $\text{WT}(\vec{x}) \leq e$, else determine that $\text{WT}(\vec{x}) > e$. (A natural question is how many tests do e -disjunct matrices require: we will come back to this question in a bit.) This seems to let us test whether \vec{y} is within a Hamming distance of e from some codeword or not. We would like to point out that the above is essentially reducing one instance of the tolerant testing problem to multiple instances of error-detection.

Thus, all we need to do is come up with a way to implement the tests to \vec{x} . A natural option, which we take, is that for any test $Q \subseteq [n]$, we check if $\vec{y}_Q \in \text{RS}_Q[k]$, where $\text{RS}_Q[k]$ is the RS code (of dimension k) projected onto Q . This immediately puts in one restriction: we will need $|Q| \geq k$ (as otherwise every test will return a positive answer). However, there is another subtle issue that makes our analysis more complicated— we do not necessarily have that $\vec{x}_Q = \vec{0}$ iff $\vec{y}_Q \in \text{RS}_Q[k]$. While it is true that $\vec{x}_Q = \vec{0}$ implies $\vec{y}_Q \in \text{RS}_Q[k]$, the other direction is not true. The latter is possible only if \vec{y} agrees with some codeword $\vec{c}' \neq \vec{c}$ in the positions indexed by Q . Now if s is the size of the smallest test and it is the case that the only codeword that agrees with \vec{y} in at least s positions is \vec{c} , then we'll be done. We show that this latter condition is true for RS codes if $s \geq e + k$ for worst-case errors or with high probability if $s \geq 4k$ for random errors.

It is now perhaps not surprising that the list decodability of the code plays a role in our general result for worst-case errors. Assume that the code C under consideration is $(n - s, L)$ list decodable (i.e. every Hamming ball of radius $n - s$ has at most L codewords in it) and one can do error detection on C projected down to any test of size at least s . If we pick our disjunct matrix carefully and L is not too large, it seems intuitive that one should be able to have, for most of the tests, that $\vec{x}_Q \neq \vec{0}$ implies $\vec{y}_Q \notin C_Q$. We are able to show that if the matrix is picked at random, then this property holds. In addition, it is the case that the “decoding” of \vec{x} from the result of the test can be done even if some of the test results are faulty (i.e. $\vec{y}_Q \in C_Q$ even though $\vec{x}_Q \neq \vec{0}$). The proof of this fact requires a fair bit of work: we will come back to the issues in a bit.

Another side-effect of the fact that our algorithm does not readily translate into the group testing scenario is that even though we have been able to salvage the case for $\text{WT}(\vec{x}) \leq e$, we can no longer guarantee that if $\text{WT}(\vec{x}) > e$, that our algorithm will catch it. In the latter case, our algorithm might say $\text{WT}(\vec{x}) > e$ or it might return a subset $S \subseteq [n]$ that purportedly contains all the error locations. However, we can always check if $\vec{y}_{[n] \setminus S} \in \text{RS}_{[n] \setminus S}[k]$ to rule out the latter case. This seems to require another pass on the input but we are able to implement the final algorithm in one pass by giving a one-pass algorithm for the following problem: Given as input \vec{y} followed by $T \subseteq [n]$

such that $|T| = e$, design a one-pass $O(e \log n)$ space algorithm to check if $\vec{y}_{[n] \setminus T} \in \text{RS}_{[n] \setminus T}[k]$. The main idea is to encode the locations in T as an unknown degree e polynomial and to fill in the unknown coefficients once the algorithm gets to T in the input.

We now return to the question of how many tests we can get away with while using e -disjunct matrices. The best known construction uses $O(e^2 \log n)$ tests [13] and this is tight to within a $\log e$ factor (cf. [15]). Thus, to get sublinear space, we need to have $e = o(\sqrt{n})$. To break the \sqrt{n} barrier, instead of e -disjunct matrices, we use the recently discovered notion of the (e, e) -list disjunct matrix [18]. An (e, e) -list disjunct matrix has the property that when applied to \vec{x} such that $\text{WT}(\vec{x}) \leq e$, it returns a subset $S \subseteq [n]$ such that (i) $x_i = 1$ implies $i \in S$ and (ii) $|S| \leq \text{WT}(\vec{x}) + e$. It is known that such matrices exist with $O(e \log n)$ rows. In Section 8, we show that such matrices can be constructed with $O(e \log^2 n)$ random bits. However, note we can now only distinguish between the cases of $\text{WT}(\vec{x}) \leq e$ and $\text{WT}(\vec{x}) \geq 2e$.

The use of list disjunct matrices also complicates our result for worst case errors that uses the list decodability of the code under consideration. The issue is that when we pick the desired matrix at random, with the extra task of “avoiding” all of the $L - 1$ codewords other than \vec{c} that can falsify the answer to the test, we can only guarantee that the “decoding” procedure is able to recover a constant fraction of the positions in error. This is similar to the notion of error reduction in [23]. This suggests a natural, iterative $O(\log e)$ -pass algorithm. Using our earlier trick, we can again implement our algorithm in one pass. Finally, the plain vanilla proof needs $\Omega(n)$ random bits. We observe that the proof goes through with limited independence and use this to reduce the amount of randomness to $O(e^2 \log^3 n)$ bits. Reducing the random bits to something smaller, such as $O(e \log n)$, is an open problem.

The speedup in the runtime from the naive $\tilde{O}(n^2)$ to $\tilde{O}(ne)$ for the tolerant testing algorithms is obtained by looking at certain explicit disjunct matrices and observing that the reduced error detection problems are nicely structured.

There are two unsatisfactory aspects of our algorithms: (i) The $O(e \log^2 n)$ space complexity and (ii) The condition that $e \leq O(n/k)$ (which in turn follows from the fact that we have $s = n/(2e)$). We show that both of these shortcomings are essentially unavoidable with our techniques. In particular, a lower bound on the 1^+ decision tree complexity of the threshold function from [5] implies that at least $\Omega(e)$ invocations of the error detection routine are needed. Further, we show that for sublinear test complexity, the support size s must be in $O(\frac{n}{e} \log n)$. This follows by interpreting the reduction as a set cover problem and observing that any set covers only a very small fraction of the universe.

2 Preliminaries

We begin with some notation. Given an integer m , we will use $[m]$ to denote the set $\{1, \dots, m\}$. We will denote by \mathbb{F}_q the finite field with q elements. An $a \times b$ matrix M over \mathbb{F}_q will be called strongly explicit if given any $(i, j) \in [a] \times [b]$, the entry $M_{i,j}$ can be computed in space $\text{poly}(\log q + \log a + \log b)$. Given a vector $\vec{y} \in \Sigma^n$ ($C \subseteq \Sigma^n$ resp.) and a subset $S \subseteq [n]$, we will use \vec{y}_S (C_S resp.) to denote \vec{y} (vectors in C resp.) projected down to the indices in S . We will use $\text{WT}(\vec{x})$ to denote the number of non-zero entries in \vec{x} . Further, for $S \subseteq [n]$, we will use $\text{WT}_S(\vec{x})$ to denote $\text{WT}(\vec{x}_S)$.

Codes. A code of *dimension* k and *block length* n over an alphabet Σ is a subset of Σ^n of size $|\Sigma|^k$. The *rate* of such a code equals k/n . A code C over \mathbb{F}_q is called a linear code if C is a linear subspace of \mathbb{F}_q^n . If C is linear, then it can be described by its parity-check matrix H , i.e. for every

$\vec{c} \in C$, $H \cdot \vec{c}^T = \vec{0}$. An asymptotically good code has constant rate and constant relative distance (i.e. any two codewords differ in at least some fixed constant fraction of positions).

Tolerant Testers. We begin with the central definition. Given a code $C \subseteq \Sigma^n$, reals $0 \leq d < c \leq 1$, $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$ and integers $r = r(n)$ and $s = s(n)$, an $(r, s, \varepsilon_1, \varepsilon_2)_{c,d}$ -tolerant tester \mathcal{T} for C is a randomized algorithm with the following properties for any input $\vec{y} \in \Sigma^n$: (1) If $\Delta(\vec{y}, C) \leq \varepsilon_1 n$, then \mathcal{T} accepts with probability at least c ; (2) If $\Delta(\vec{y}, C) \geq \varepsilon_2 n$, then \mathcal{T} accepts with probability at most d ; (3) \mathcal{T} makes at most r passes over \vec{y} ; and (4) \mathcal{T} uses at most s space for its computation.

Further, we will consider the following special cases of an $(r, s, \varepsilon_1, \varepsilon_2)_{c,d}$ -tolerant tester: (i) An $(r, s, 0, \varepsilon)_{c,d}$ -tolerant tester will be called an $(r, s, \varepsilon)_{c,d}$ -local tester. (ii) An $(r, s, 0, 1/n)_{c,d}$ -tolerant tester will be called an $(r, s)_{c,d}$ -error detector. There are some definitional issues that are resolved in Section 4.

List Disjunct Matrices. We give a low-space algorithm that can compute a small set of possible defectives given an outcome vector which is generated by a list disjunct matrix. Relevant definitions and material related to the algorithm can be found in Section 6.

Some Explicit Families of Codes. We now mention two explicit families of codes that we will see later on in the paper. We first begin with the Reed-Solomon code. Given $q \geq n \geq 1$ and a subset $S = \{\alpha_1, \dots, \alpha_n\} \subseteq \mathbb{F}_q$, the Reed-Solomon code with *evaluation set* S and dimension k , denoted by $\text{RS}_S[k]$, is defined as follows: Any message in \mathbb{F}_q^k naturally defines a polynomial $P(X)$ of degree at most $k-1$ over \mathbb{F}_q . The codeword corresponding to the message is obtained by evaluating $P(X)$ over all elements in S . It is known that a $(n-k) \times n$ parity check matrix of RS_S is given by $H_{\text{RS}_S} = \{v_j \cdot \alpha_j^i\}_{i=0}^{n-k-1}, j=1, \dots, n$, where

$$v_j = \frac{1}{\prod_{1 \leq \ell \leq n, \ell \neq j} (\alpha_j - \alpha_\ell)}$$

Another explicit code family we will consider are expander codes. These are binary codes whose parity check matrices are incidence matrices of constant-degree bipartite expanders. In particular, if we start with a strongly explicit expander, then the parity check matrix of the corresponding expander code will also be strongly explicit.

3 Connections to Property Testing

3.1 The basic connection

We now highlight a simple connection between tolerant testers in the data stream world and tolerant testers in the query world:

Proposition 1. *Let $C \subseteq [q]^n$ be such that it has a tolerant tester \mathcal{T} with query complexity r , thresholds ε_1 and ε_2 and time complexity $t_q(r)$ (i.e. it makes $t_q(r)$ operations over $[q]$ for any possible query realization). Then there also exists an $(r, O(t_q(r) + r \log n), \varepsilon_1, \varepsilon_2)_{c,s}$ -tolerant tester \mathcal{T}' . Further, if \mathcal{T} is non-adaptive, then \mathcal{T}' can be implemented as a $(1, O(t_q(r) + r \log n), \varepsilon_1, \varepsilon_2)_{c,s}$ -tolerant tester.*

Proof. The claimed result follows from the obvious simulation. In general, the tester \mathcal{T}' works as follows: \mathcal{T} queries r positions in the input and then applies some function on the queried values

(in the case when \mathcal{T}' is adaptive, it applies possibly r different functions after each query). As the total time complexity of \mathcal{T} is $t_q(r)$, the entire computation of \mathcal{T} can be done in time (and hence, space) $t_q(r)$. If \mathcal{T} is non-adaptive, all the query positions can be decided upfront and all the values can be determined in one pass. Otherwise the simulation might need r passes in the worst case. We might need to use an additional $O(r \log n)$ space to store the indices of the query positions, which implies that \mathcal{T}' has the claimed properties. ■

Remark 1. In general, one cannot say much about $t_q(r)$ other than bounding it by $2^{O(q^r)}$ as the definition of the usual tolerant tester does not put any computational efficient constraints on the testers. However, if the tolerant tester \mathcal{T} makes a constant number of queries then its time complexity is also a constant number of operations over the alphabet.

3.2 Tolerant Testing of Reed-Muller Codes

Let $\text{RM}(q, \ell, m)$ denote the Reed-Muller code obtained by evaluating m -variate polynomials over \mathbb{F}_q of total degree $\ell < q$. These codes are known to have block length $n = q^m$, dimension $\binom{m+\ell}{m}$ and distance $(1 - \ell/q)n$ (cf. [24, Lect. 4]). Note that this implies that if m is a constant and $\ell = \Omega(q)$, then $\text{RM}(q, \ell, m)$ is asymptotically good. These codes are known to be tolerant testable in the property testing world with polynomial number of queries.

Theorem 2 ([16]). *Let $m, \ell, q \geq 1$ be integer such that $\ell < c \cdot q$ for some universal constant c . Then there exists a tolerant tester for $\text{RM}(q, \ell, m)$ in the property testing world that can distinguish between at most $\varepsilon_1 n$ and at least $\varepsilon_2 n$ errors with $q = n^{1/m}$ queries, where ε_1 and ε_2 are absolute constants that only depends on c .*

In fact, the test is simple to describe: pick a random line in \mathbb{F}_{q^m} and check if the projected down received word is within some threshold Hamming distance from the corresponding RS code of dimension ℓ and block length q . This latter step can be solved using the fast list decoding algorithm for RS codes from [2] in time $O(q \log^2 q)$ (and hence, in the same amount of space). Thus, Proposition 1 implies the existence of an asymptotically good code that can be tolerant tested by a one-pass, $O(n^\varepsilon)$ space (for any $\varepsilon > 0$) data stream algorithm. (Note that the code has an absolute constant as its relative distance but its rate is exponentially small in $1/\varepsilon$.)

4 Some definitional issues

One decision that we need to make is how we count the space/time requirement for our algorithms. In particular, given a code defined over Σ , do we do our accounting in terms of number of operations over Σ or the number of operations over “bits”? This question is moot when Σ has constant size as both the measures will be within constant factors of each other. However, if $|\Sigma|$ can depend on n , which will be the case in some of the codes that we consider in this paper, the two measures will not be within constant factors anymore. In particular, for arbitrary Σ , an operation over Σ may take $\Omega(n)$ space and time, which will be prohibitive for our purposes.

We resolve the question above in the following way: First, we will account for the complexity measures in terms of the number of operations in Σ . Further, for positive results, we will focus on the case where $\Sigma = \mathbb{F}_q$ with $q \leq n^{O(1)}$. Note that in this case, all operations (including addition, multiplication and exponentiation) can be done in $\text{poly}(\log n)$ time and $O(\log n)$ space. Finally, for

general fields \mathbb{F}_q , the algorithm will also need access to an irreducible polynomial (of degree at most $O(\log n)$). However, note that the definition of a code provides the definition of its alphabet. We will assume that the algorithm has full prior knowledge about the code (including e.g. the value of n). Thus, we will assume that the irreducible polynomial will be (implicitly) a part of the input to the algorithm. For certain cases, when the irreducible polynomial is part of an explicit family, the algorithms can compute these irreducible polynomials “on the fly” and thus, do not need to be part of the input.

Finally, by definition, the block length of a code is fixed. However, for a meaningful asymptotic analysis, we need to think of an increasing sequence of block lengths. Thus, from now on when we talk about a code, we implicitly mean a family of codes.

5 Data Stream Algorithms for Error-Detection

A positive result. We first show that any linear code with a strongly explicit parity check matrix has an efficient 1-pass data stream error detector.

Note that for a linear code $C \subseteq \mathbb{F}_q^n$ with parity check matrix H , the error detection problem with the usual polynomial time complexity setting is trivial. This is because by the definition of parity check matrix for any $\vec{y} \in \mathbb{F}_q^n$, $\vec{y} \in C$ if and only if $H \cdot \vec{y}^T = \vec{0}$. However, the naive implementation requires $\Omega(n)$ space which is prohibitive for data stream algorithms. We will show later that for deterministic data stream algorithms with a constant number of passes, this space requirement is unavoidable for asymptotically good codes.

However, the story is completely different for randomized algorithms. If we are given the *syndrome* $\vec{s} = H\vec{y}^T$ instead of \vec{y} as the input, then we just have to solve the *set equality* problem which has a very well-known one-pass $O(\log n)$ -space data stream algorithm based on the *fingerprinting* method. Because \vec{s} is a fixed linear combination of \vec{y} (as H is known), we can use the fingerprinting technique in our case. Further, unlike the usual fingerprinting method, which requires any large enough field, we need to use an extension field of \mathbb{F}_q . For this, we need to get our hands on irreducible polynomials over \mathbb{F}_q .

5.1 Families of Irreducible Polynomials

In our error detection algorithm we need low space construction of families of irreducible polynomials. Since our final algorithm will be randomized, a randomized algorithm to construct irreducible polynomials works. The following result is well-known (cf. [22, Chap. 20]):

Theorem 3. *Let q be a prime power, d be an integer and $0 < \delta < 1$ be a real number. Then there exists a randomized algorithm that outputs an irreducible polynomial of degree d over \mathbb{F}_q with probability at least $1 - \delta$. Further, this algorithm makes $O(d^4 \log(1/\delta) \log q)$ operations over \mathbb{F}_q and needs $O(\log(1/\delta) + d \log q)$ bits of space.*

Coming up with a deterministic polynomial algorithm for construction of irreducible polynomials is an open question. However, it turns out that in our application, we would be happy if the final irreducible polynomial has degree d' such that $d' \geq d$ and is not much larger than d . In particular for *prime* p , there exists a deterministic algorithm that runs in time (and hence, space) $\text{poly}(d \log p)$ and outputs a polynomial with degree at least d and at most $O(d \log p)$ [1].

Next, we show that constructing such irreducible polynomials can also be done for fields of characteristic 2. The result follows from other known results.

Theorem 4. *Let q be a power of 2 and let $d \geq 1$ be an integer. Given the irreducible polynomial that generates \mathbb{F}_q , there exists a deterministic $O(d \log q)$ space, $O((d^2 + \log q) \log^2 q)$ -time algorithm that computes an irreducible polynomial over \mathbb{F}_q with degree d' such that $d \leq d' \leq 2d$.*

Coming up with an analogous result to Theorem 4 for odd characteristic seems to be an open problem.

We begin the proof of Theorem 4. We will use the following result:

Theorem 5 (cf. [9]). *Let $m \geq 1$ be an integer and let $\beta \in \mathbb{F}_{2^m}$ such that $\text{Tr}(\beta) \neq 0$, where $\text{Tr}(x) = \sum_{i=0}^{m-1} x^{2^i}$ is the trace function. Define the polynomials $A_k(X)$ and $B_k(X)$ recursively as follows (for $k \geq 0$):*

$$\begin{aligned} A_0(X) &= X \\ B_0(X) &= 1 \\ A_{k+1}(X) &= A_k(X)B_k(X) \\ B_{k+1} &= A_k^2(X) + B_k^2(X). \end{aligned}$$

Then $A_k(X) + \beta \cdot B_k(X)$ is an irreducible polynomial over \mathbb{F}_{2^m} of degree 2^k .

To begin with, let us assume we can get our hands on a β as required in Theorem 5. Given such a β , the rest of the proof is simple. Pick k to be the smallest integer such that $d' = 2^k \geq d$. It is easy to check that $d \leq d' \leq 2d$ as required. To compute the final irreducible polynomial, we will need to do k iterations to compute $A_i(X)$ and $B_i(X)$ (for $1 \leq i \leq k$). It is easy to check that each iteration requires $O(2^i \log q)$ space (to store the intermediate polynomials) and $O(2^{2i} \log^2 q)$ time (to compute the product of two polynomials of degree at most 2^i). To complete the proof of Theorem 4, we show how to efficiently compute an appropriate β .

We claim that β can be chosen to be α^i for some $0 \leq i \leq m-1$, where we use $\{1, \alpha, \dots, \alpha^{m-1}\}$ as the standard basis for \mathbb{F}_{2^m} , for some root α of the irreducible polynomial that generates \mathbb{F}_{2^m} .² To see why this is true, assume for the sake of contradiction that $\text{Tr}(\alpha^i) = 0$ for every $0 \leq i \leq m-1$. Then as every $\gamma \in \mathbb{F}_{2^m}$ can be written as a linear combination of $1, \alpha, \dots, \alpha^{m-1}$, $\text{Tr}(\gamma) = 0$ (this follows from the well-known fact that $\text{Tr}(\gamma_1 + \gamma_2) = \text{Tr}(\gamma_1) + \text{Tr}(\gamma_2)$). This implies that $\text{Tr}(X)$ has 2^m roots, which is a contradiction as $\text{Tr}(X)$ is a non-zero polynomial of degree 2^{m-1} . Finally, the correct choice of $\beta = \alpha^i$ can be determined by going through all $0 \leq i \leq m-1$ and evaluating $\text{Tr}(\alpha^i)$ (which can be done in $O(\log^3 q)$ time).

We now state our result.

Theorem 6. *Let $C \subseteq \mathbb{F}_q^n$ be a linear code of dimension k and block length n with parity check matrix $H = \{h_{i,j}\}_{i=0}^{n-k-1, n}$. Further, assume that any entry $h_{i,j}$ can be computed in space $\mathcal{S}(n, q)$, for some function \mathcal{S} . Given an $a \geq 1$, there exists a $(1, O(\mathcal{S}(n, q) + a \log n))_{1, n-a}$ -error detector for C .*

² Note that if m is odd, then just $\beta = 1$ suffices.

5.2 Proof of Theorem 6

Let $\vec{y} = (y_1, \dots, y_n) \in \mathbb{F}_q^n$ be the received word and let $\vec{s} = (s_0, \dots, s_{n-k-1}) = H\vec{y}^T \in \mathbb{F}_q^{n-k}$. It is easy to check that for every $0 \leq i < n-k$,

$$s_i = \sum_{j=1}^n y_j h_{i,j}. \quad (1)$$

Further define

$$S(X) = \sum_{i=0}^{n-k-1} s_i X^i.$$

Note that our task is to verify whether $S(X)$ is the all zeros polynomial. Towards this end, we will use the fingerprinting technique.

Let $Q = q^d$ for some d to be chosen later such that $n^{1+\alpha} \leq Q \leq qn^{1+\alpha}$. The algorithm is simple: pick a random $\beta \in \mathbb{F}_Q$ and verify if $S(\beta) = \sum_{i=0}^{n-k-1} s_i \beta^i = 0$. By (1), this is the same as checking if $\sum_{i=0}^{n-k-1} \left(\sum_{j=1}^n y_j h_{i,j} \right) \beta^i = 0$. Note that this is possible as \mathbb{F}_Q is an extension field of \mathbb{F}_q and thus, all the terms in the sum belong to \mathbb{F}_Q . Thus, by changing the order of sums, we need to verify if

$$\sum_{j=1}^n y_j \left(\sum_{i=0}^{n-k-1} \beta^i h_{i,j} \right) = 0. \quad (2)$$

It is easy to verify that the above sum can be computed in one pass as long as the quantity $\sum_{i=0}^{n-k-1} \beta^i h_{i,j}$ can be computed efficiently “on the fly.” The latter is possible as we know β and we can compute any entry $h_{i,j}$ on the fly. If $S(X)$ is the all zeros polynomial, then the check will always pass. If on the other hand, $S(X)$ is a non-zero polynomial of degree at most n , then $S(\beta) = 0$ for at most n values $\beta \in \mathbb{F}_Q$. Thus, the probability of the check passing is at most n/Q which by our choice of Q is at most $1/n^\alpha$.

To complete the proof we need to analyze the space requirement of the algorithm above. First we note that by Theorem 3 we can compute an irreducible polynomial over \mathbb{F}_q of degree $d = \left\lceil 2 \frac{\log n}{\log q} \right\rceil$. Note that $n^{1+\alpha} \leq Q \leq qn^{1+\alpha}$ as claimed before. Also note that any operation in \mathbb{F}_Q can be carried out by storing $O(d)$ elements from \mathbb{F}_q . This implies that the sum in (2) can be computed with space $O(\mathcal{S}(n, q) + d) = O(\mathcal{S}(n, q) + \alpha \log n)$.

An inspection of the proof above shows that the time complexity of the error detector is dominated by the number of \mathbb{F}_q operations needed to compute $\sum_{i=0}^{n-k-1} \beta^i h_{i,j}$.

For expander codes, this time complexity is just a constant number of \mathbb{F}_q operations (and hence $O(\log n / \log q)$ operations in \mathbb{F}_q). For RS codes, recall that we have $h_{i,j} = v_j \cdot \alpha_j^i$, where $v_j = \frac{1}{\prod_{1 \leq \ell \leq n, \ell \neq j} (\alpha_j - \alpha_\ell)}$. If, say, for some fixed $\beta \in \mathbb{F}_q^*$, $v_j = \beta$ for every $1 \leq j \leq n$, then one can essentially ignore v_j and one only needs to compute $\sum_{i=0}^{n-k-1} \beta^i \alpha_j^i$, which is just $\frac{(\beta \alpha_j)^{n-k} - 1}{\beta \alpha_j - 1}$ unless $\beta = 0$ (in which case the sum is 0) or $\beta = (\alpha_j)^{-1}$ (in which case the sum is just $(n-k)$ modulo the characteristic of \mathbb{F}_q). The latter condition can be verified with $\text{poly}(\log n / \log q)$ operations in \mathbb{F}_q .

In general RS, any $h_{i,j}$ can be computed with $\tilde{O}(n)$ operations in \mathbb{F}_q . Thus, the sum can be computed in time $\tilde{O}(n^2)$.

Thus, we have argued that

Corollary 7. *Let q be a prime power and define $S = \{\alpha_1, \dots, \alpha_n\}$. Then there exists an $(1, O(\log n))_{1,1/2}$ -error detector for RS_S that runs in time $\tilde{O}(n^2)$. Further, if there exists a $\beta \in \mathbb{F}_q^*$ such that for every $1 \leq j \leq n$, $\prod_{1 \leq \ell \leq n, \ell \neq j} (\alpha_j - \alpha_\ell) = \beta$, then the algorithm can be implemented in $\tilde{O}(n)$ time.*

It is easy to check that $\mathcal{S}(q, n)$ is $O(\log n)$ for (strongly explicit) expander codes and RS (and GRS) codes. This implies the following:

Corollary 8. *Let $n \geq 1$. Then for $q = 2$ and $n \leq q \leq \text{poly}(n)$, there exists an asymptotically good code $C \subseteq \mathbb{F}_q^n$ that has a $(1, O(\log n))_{1,1/2}$ -error detector.*

A negative result. We show that randomness is necessary even for local testing. In particular, we show the following:

Theorem 9. *Let $C \subseteq [q]^n$ be a code of rate R and relative distance δ and let $0 \leq \varepsilon \leq \delta^2/8$ be a real number. Then any $(r, s, \varepsilon)_{1,0}$ -local tester for C needs to satisfy $r \cdot s \geq \frac{\delta R n}{6}$.*

5.3 Proof of Theorem 9

We will be using communication complexity to prove Theorem 9.

The proof uses the standard fooling set technique, however, unlike the usual application of two-party communication complexity in data stream algorithms, where the stream is broken up into two fixed portions, in our case we need to be careful about how we divide up the input. To see the necessity of this, consider the code $C \times C \subseteq \Sigma^{2n}$ and say we break the received word \vec{y} in the middle and assign the first half (call it \vec{y}_1) to Alice and the second half to Bob. In this case there is a simple $O(\log n)$ protocol—Alice simply sends the distance of \vec{y}_1 to the closest codeword in C to Bob—to compute the distance of \vec{y} to the closest codeword in $C \times C$ *exactly*. However, we can show that for every asymptotically good code, there is some way of breaking up the input into two parts such that there exists an exponentially sized fooling set.

To further explain, we do a quick recap of some of the basic concepts in communication complexity and refer the reader to source material for more details [19].

Let $g : \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \rightarrow \{0, 1\}$ be a function. Further assume Alice has $x \in \{0, 1\}^{n_1}$ and Bob has $y \in \{0, 1\}^{n_2}$. The (deterministic) communication complexity of g , denoted by $\text{CC}(g)$, is the minimum number of bits that Alice and Bob must exchange in order to determine $g(x, y)$ in the worst case. The following observation is a standard technique to obtain lower bounds for data stream algorithms:

Proposition 10. *Let \mathcal{A} be an r -pass, s -space deterministic data stream algorithm that decides g . Then $r \cdot s \geq \text{CC}(g)$.*

Next we consider the following technique for lower bounding the communication complexity of a function. A subset $F \subseteq \{0, 1\}^{n_1} \times \{0, 1\}^{n_2}$ is called a *fooling set* for g if (i) For every $(x, y) \in F$, $g(x, y) = b$ for some fixed $b \in \{0, 1\}$ and (ii) For every $(x_1, y_1) \neq (x_2, y_2) \in F$, either $g(x_1, y_2) = 1 - b$ or $g(x_2, y_1) = 1 - b$. The following result is well-known:

Proposition 11 (cf. [19]). *Let $g : \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \rightarrow \{0, 1\}$ and F be a fooling set for g . Then $\text{CC}(g) \geq \log(|F|)$.*

Finally, we will consider boolean functions with one input and we define their communication complexity as follows: Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$. Further, for any $0 \leq n_1, n_2 \leq n$ such that $n_1 + n_2 = n$, define $f_{n_1, n_2} : \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \rightarrow \{0, 1\}$ by naturally “dividing” up the n -bit input for f into the two required inputs for f_{n_1, n_2} . The communication complexity of f is then defined as follows:

$$\text{CC}(f) = \max_{\substack{0 \leq n_1, n_2 \leq n, \\ n_1 + n_2 = n}} \text{CC}(f_{n_1, n_2}).$$

We are now ready to prove Theorem 9. We will do so by proving that the deterministic communication complexity of the following function is large. Define $f_C : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $f_C(\vec{y}) = 1$ if $\vec{y} \in C$; $f_C(\vec{y}) = 0$ if $\Delta(\vec{y}, C) \geq \delta^2 n/8$; otherwise $f_C(\vec{y})$ can take an arbitrary value. We will show that:

Lemma 12. *f_C has a fooling set of size at least $q^{\delta Rn/6}$.*

Note that as we are measuring space in terms of the number of elements from $[q]$, Lemma 12 and Proposition 11 imply Theorem 9.

In the rest of the section, we prove Lemma 12. For notational convenience, define $k = Rn$, $d = \delta n$, $\alpha = \frac{\delta}{4}$ and $\beta = \frac{\delta}{6}$. Thus we need to exhibit a fooling set of size at least $q^{\beta k}$.

Our fooling set will be a subset $F \subseteq C$ with an $0 < n_1 < n$ such that $F_{[n_1]}$ and $F_{[n] \setminus [n_1]}$ have distance at least $\alpha d/2$. Further for any $\vec{c} \in T$ and $\vec{c}' \in C \setminus F$, $\Delta(\vec{c}, \vec{c}') \geq \alpha d/2$. We claim that such an F is indeed a fooling set. To see this consider $\vec{c}^1 \neq \vec{c}^2 \in F$, where $\vec{c}^1 = (c_1^1, c_2^1)$, $\vec{c}^2 = (c_1^2, c_2^2)$, $c_1^i \in F_{[n_1]}$ and $c_2^i \in F_{[n] \setminus [n_1]}$. By definition, $f_C(\vec{c}^1) = f_C(\vec{c}^2) = 1$. Next we show that either $f_C(\vec{y}_1) = 0$ or $f_C(\vec{y}_2) = 0$, where $\vec{y}_1 = (c_1^1, c_2^2)$ and $\vec{y}_2 = (c_1^2, c_2^1)$. For any $\vec{c} \in C \setminus F$, this is true by definition of F . For any $\vec{c} \in F$, this is true by the distance properties of $F_{[n_1]}$ and $F_{[n] \setminus [n_1]}$.

We will construct the fooling set F in a greedy fashion. We begin with the case when $n < 2(1 - \alpha)d$. We claim that in this case $F = C$ and $n_1 = \lfloor n/2 \rfloor$ works. Note that both n_1 and $n - n_1$ are both at most $(1 - \alpha)d$. Since C has distance d , this implies that both $F_{[n_1]}$ and $F_{[n] \setminus [n_1]}$ have distance at least $d - (1 - \alpha)d = \alpha d$. This completes the proof for the base case.

For the general $n \geq 2(1 - \alpha)d$ case, we reduce it to the base case. In particular, we present a greedy iterative process, where at the end of the i^{th} step, we have a subset $F_i \subseteq C$, with the property that for every $\vec{c} \in F_i$ and $\vec{c}' \in C \setminus F_i$, $\Delta(\vec{c}, \vec{c}') \geq \alpha d/2$. We of course start with $F_0 = C$. It will turn out that we will run this process for $r \leq \frac{1}{2\alpha}$ times and $F = F_r$. For ease of exposition, we will also track variables m_i and d_i such that $m_0 = n$ and $d_0 = d$. Next, we mention the invariance that we will maintain with the iterative algorithm. First, it will always be the case that $m_{i+1} = m_i - (1 - \alpha)d_i$ and $d_{i+1} \geq d_i - \alpha d_i$. Think of m_i as the block length of the (projected down) F_i and d_i as the corresponding distance.

Assume that we have our hands on F_i . If $m_i < 2(1 - \alpha)d_i$, then we are in the base case and the process terminates. (In this case $r = i$ and $n_1 = \lfloor m_r/2 \rfloor$.) If not, then F_i projected onto the first m_i positions (call this projected down code G_i) has distance at least d_i . Group codewords in F_i such that in each cluster the codewords in G_i projected down to the last $(1 - \alpha)d_i$ positions differ from each other in $< \alpha d_i$ positions. If the number of clusters is at least $q^{\beta k}$, then let F_{i+1} be defined by picking one codeword from each of the clusters and the process terminates. (In this case $r = i + 1$ and $n_1 = m_r$.) If not, then define F_{i+1} to be the largest cluster and define $m_{i+1} = m_i - (1 - \alpha)d_i$. We then continue the process for $i + 1$.

For the time being, assume that the following are true: (i) The process stops at iteration r such that $r \leq \frac{1}{2\alpha}$; (ii) G_i has distance d_i such that $d_i \geq (1 - \alpha)d$; and (iii) Every codeword in F_i differs

from every codeword in $F_{i-1} \setminus F_i$ in at least $\alpha d/2$ positions. Assuming these three properties, we argue that F_r indeed has the required properties.

Assume that the process terminates when the base case is reached. Let $G = G_r$. Then by the argument in the base case, we have that both $G_{[n_1]}$ and $G_{[m_r] \setminus [n_1]}$ (and hence, $F_{[n_1]}$ and $F_{[n] \setminus [n_1]}$) have distance at least $\alpha d_r \geq \alpha d/2$, where the inequality follows from properties (i) and (ii). Further, by property (iii), it is the case that for every $\vec{c} \in F$ and $\vec{c}' \in C \setminus F$, $\Delta(\vec{c}, \vec{c}') \geq \alpha d/2$. Finally, note that when we pick a single cluster, we have $|F_{i+1}| \geq |F_i|/q^{\beta k}$. Thus, if we terminate with the base case, we have

$$|F| \geq \frac{q^k}{q^{\beta k}} = q^{k(1-r\beta)} \geq q^{\beta k},$$

where the last inequality follows from the following argument for the inequality $1 - r\beta \geq \beta$. This inequality is satisfied if

$$r \leq \frac{1}{\beta} - 1.$$

Now by property (i), $r \leq \frac{1}{2\alpha} = \frac{2}{\delta} = \frac{3}{\beta}$. Now as $\beta \leq \delta/6 \leq 1/6$, we have $\beta \leq 1/(3\beta) \leq 1/\beta - 1$, as desired.

Now we consider the case when the process terminates before reaching the base case. In this case because of the termination condition, we have that G_{r-1} projected onto the last $(1-\alpha)d_{r-1}$ positions (and hence, $F_{[n] \setminus [n_1]}$) has distance at least $\alpha d_{r-1} \geq \alpha d/2$, where the inequality follows from properties (i) and (ii). Also $F_{[n_1]}$ has distance at least $d_{r-1} - (1-\alpha)d_{r-1} = \alpha d_{r-1} \geq \alpha d/2$. Further, by property (iii), it is the case that for every $\vec{c} \in F$ and $\vec{c}' \in C \setminus F$, $\Delta(\vec{c}, \vec{c}') \geq \alpha d/2$. Finally, by the termination condition, we have $|F| \geq q^{\beta k}$, as desired.

Thus, we are done with the proof modulo showing that properties (i)-(iii) hold, which is what we do next. We begin with property (i). Note that if we do not terminate in the middle, then we have $d_{i+1} \geq (1-\alpha)d_i \geq (1-\alpha)^i d$. Now note that

$$m_r = n - \sum_{i=0}^{r-1} (1-\alpha)d_i \leq n - d \sum_{i=1}^r (1-\alpha) = n - \frac{(1-\alpha)(1-(1-\alpha)^r)d}{\alpha}.$$

Since $m_r \geq 0$, we have

$$1 - (1-\alpha)^r \leq \frac{\alpha n}{(1-\alpha)d} = \frac{\alpha}{(1-\alpha)\delta} \leq 1 - \exp(-1/2),$$

where the last inequality follows from the fact that $\alpha = \delta/4 \leq 1/4$. Thus the above implies that

$$(1-\alpha)^r \geq \exp(-1/2),$$

which in turn implies

$$r \ln \left(\frac{1}{1-\alpha} \right) \leq \frac{1}{2}.$$

Using the fact that $\ln(1-x) = -(x + x^2/2 + x^3/3 + \dots)$ for $|x| < 1$, we get that the above implies

$$r(\alpha + \alpha^2/2 + \alpha^3/3 + \dots) \leq \frac{1}{2},$$

which in turn implies that $r\alpha \leq \frac{1}{2}$, as desired.

We now move to property (ii). As we saw earlier, we have $d_i \geq (1 - \alpha)^i d$ for $i \leq r$. Now as $\alpha r \leq 1/2$ (and hence $\alpha i \leq 1/2$), we have that $(1 - \alpha)^i \geq 1 - i\alpha$, which proves property (ii). Note that this also implies that $d_i \geq d/2$. Finally, for property (iii), note that by construction, if we do not terminate in middle at step i , every codeword in F_{i+1} differs from $F_i \setminus F_{i+1}$ in at least αd_i positions. Since, $d_i \geq d/2$, property (iii) follows. The proof is complete.

Error detection of a projected down code. We will be dealing with $RS_S[k]$ with $S = \{\alpha_1, \dots, \alpha_n\}$. In particular, we are interested in a one-pass, low space data stream algorithm to solve the following problem: The input is $\vec{y} \in \mathbb{F}_q^n$ followed by a subset $E \subseteq S$ with $|E| = e$. We need to figure out if $\vec{y}_{S \setminus E} \in RS_{S \setminus E}[k]$. We have the following result:

Lemma 13. *Let $e, n, k \geq 1$ be integers such that $k + e \leq n$. Then the problem above can be solved by a one-pass, $O(e + a \log n)$ space data stream algorithm with probability at least $1 - n^{-a}$, for any $a \geq 1$.*

5.4 Proof of Lemma 13

Consider the degree e polynomial $P_E(X) = \prod_{i \in E} (X - \alpha_i)$. Further, consider a new received word $\vec{z} = (z_1, \dots, z_n)$ where $z_i = y_i \cdot P_E(\alpha_i)$. The algorithm to solve the problem above just checks to see if $\vec{z} \in RS_S[e + k]$.

We begin with the correctness of the algorithm above. If $\vec{y}_{[n] \setminus E} \in RS_{S \setminus E}[k]$, that is, $\vec{y}_{[n] \setminus E}$ is the evaluation of a polynomial $f(X)$ of degree at most $k - 1$ over $S \setminus E$, then \vec{z} is the evaluation of $f(X) \cdot P_E(X)$ over S . In other words, $\vec{z} \in RS_S[k + e]$.

Now it turns out that the other direction is also true. That is, if $\vec{z} \in RS_S[k + e]$ then $\vec{y}_{[n] \setminus E} \in RS_{[n] \setminus E}[k]$. Note that \vec{z} is the evaluation of a degree at most $e + k - 1$ polynomial $g(X)$ over S , where $g(X) = P_E(X) \cdot h(X)$, where $h(X)$ has degree at most $k - 1$. This is easy to see: by definition $P_E(X) | g(X)$ and the degree requirement on $g(X)$ implies that $h(X)$ has degree at most $k - 1$. Finally, as $P_E(\alpha_i) \neq 0$ for $i \notin E$, this implies that $h(\alpha_i) = y_i$ for $i \notin E$. In other words, $\vec{y}_{[n] \setminus E} \in RS_{[n] \setminus E}$.

We conclude this proof by showing how to deal with the unknown E using a data stream algorithm. Since E is unknown, let us denote $P_E(X) = X^e + \sum_{i=0}^{e-1} p_i X^i$, where $\{p_i\}$ are the unknown coefficients. Recall that in our error detection algorithm to check if $\vec{z} \in RS_S[e + k]$ we need to check if the following sum is 0:

$$\sum_{j=1}^n y_j P_E(\alpha_j) \left(\sum_{i=0}^{n-k-e} \beta^i h_{i,j} \right),$$

where β is a random element in a large enough extension field of \mathbb{F}_q and $\{h_{i,j}\}$ is the parity check matrix of $RS_S[k]$. Note that as $P_E(X) = X^e + \sum_{i=0}^{e-1} p_i X^i$, the sum above can be written as $Q_e + \sum_{b=0}^{e-1} p_b Q_b$, where (for $0 \leq b \leq e$)

$$Q_b = \sum_{j=1}^n y_j \alpha_j^b \left(\sum_{i=0}^{n-k-e} \beta^i h_{i,j} \right).$$

Note that each of the Q_b sums can be computed in one pass and low space without the knowledge of E .

Now the algorithm to check if $\vec{z} \in RS_S[k+e]$ is clear: maintain the $e+1$ sums \hat{Q}_b . At the end of the pass, the previous algorithm knows the set $E \subseteq [n]$. Given this, we can compute the coefficients $\{\hat{e}_b\}_{b=0}^{e-1}$. Then we declare $\vec{y}_{[n] \setminus E} \in RS_{[n] \setminus E}$ if and only if $\hat{Q}_t + \sum_{b=0}^{t-1} \hat{e}_b \hat{Q}_b = 0$. This extra computation will need storage for $O(e)$ elements in the extension field of \mathbb{F}_q .

6 Tolerant testing

In this section we assume that we are working with $RS_S[k]$, where $S = \{\alpha_1, \dots, \alpha_n\} \subseteq \mathbb{F}_q$. (However, our results will also hold for closely related codes such as the folded RS code [17].)

As was mentioned in the Introduction, there is a trivial reduction from one tolerant testing instance (say where we are interested in at most e vs. $> e$ errors) to $\binom{n}{e}$ instances of error detection: for each of the $\binom{n}{e}$ potential error locations, project the received word outside of those indices and check to see if it's a codeword in the corresponding RS code via the algorithm in Theorem 6. Using Theorem 6 (with $a = O(e)$), we can implement this as an $(1, O(e \log n), e/n, (e+1)/n)_{1, n-O(e)}$ -tolerant tester. Unfortunately, this algorithm uses $\frac{n}{e} O(e)$ time. Next, we show how to obtain roughly the same space complexity but with a much better time complexity.

Theorem 14. *Let $e, k, n \geq 1$ be integers such that $k \leq n$ and $e \leq n - k$. Then*

- (a) *If $e(e+k) \leq O(n)$, then there exists a $(1, O(e \log^2 n), e/n, 2e/n)_{1, n-O(1)}$ -tolerant tester for $RS_S[k]$ under worst-case errors.*
- (b) *If $ek \leq O(n)$, then there exists a $(1, O(e \log^2 n), e/n, 2e/n)_{1, n-O(1)}$ -tolerant tester for $RS_S[k]$ under random errors.*
- (c) *If $e \leq O(\sqrt[3]{sn/k})$, then there exists a $(1, O(e^2 \log^3 n), e/n, 5e/n)_{1, n-O(1)}$ -tolerant tester for the folded RS code with folding parameter s under worst-case errors.*

Further, all the algorithms can be implemented in $\tilde{O}(n^2)$ time.

In the above, the soundness parameter follows by picking a to be large enough while using Theorem 6. We observe that a naive implementation achieves the $\tilde{O}(n^2)$ runtime. We also show that for part (a) and (b) by replacing n by $n/\log n$ in the RHS of the upper bound on k and bumping up the space to $O(e^2 \log^2 n)$, the algorithms for $RS_{\mathbb{F}_q}[k]$ can be implemented in $\tilde{O}(ne)$ time. In fact, along with the faster running time, we get $(1, O(e \log^2 n), e/n, (e+1)/n)_{1, n-O(1)}$ -tolerant testers.

We start with some notation. Given an $t \times n$ (list) disjunct matrix M let s and s' denote the minimum and maximum Hamming weight of any row in M . Further, let $D(N)$ denote the runtime of error detector for $RS_{\alpha_1, \dots, \alpha_N}[k]$.

We begin by analyzing the runtime of the tolerant tester from part (a) of Theorem 14. The runtime has two parts: one is the time taken to run the error detector for all the projected down codes, which are determined by the rows of the (e, e) -list disjunct matrix M . Note that this step takes time at most $t \cdot D(s')$. The second part is the time taken to run the algorithm from Lemma 13, which can be verified to be $D(n)$. Thus, the overall running time is

$$t \cdot D(s') + D(n). \quad (3)$$

It can be verified that for list disjunct matrices from Section 8, $t = O(e \log n)$ and both s and s' are $\Theta(n/e)$. Further, by Corollary 7, we upper bound $D(N)$ by $\tilde{O}(N^2)$. Thus, (3) implies that the runtime is upper bounded by $\tilde{O}(n^2)$.

Next, we look at part (a) when $S = \mathbb{F}_q$. In this case we will pick an *explicit* disjunct matrix. This classic matrix is defined by associating the columns with the codewords of $\text{RS}_{\mathbb{F}_q}[k']$ for appropriate choices of q' and k' . (Note that we then have $q = n = (q')^{k'}$.) The columns of the matrix are the corresponding RS codewords, where each symbol from $\mathbb{F}_{q'}$ in the codeword is replaced by the binary vector from $\{0, 1\}^{q'}$, which has a 1 only in the position corresponding to the symbol (when thought of as an element from $[q']$). It is well known that if one picks $k' = q'/e$, then the matrix is e -disjunct and $t = O(e^2 \log^2 n)$ and $s = s' = n/\sqrt{t}$ [13]. Note that one can index the rows of this matrix by the tuples $(a, b) \in (\mathbb{F}_{q'})^2$. For the rest of the argument fix such a row (a, b) . The columns that participate in this row correspond to the messages $(m_0, \dots, m_{k'-1}) \in \mathbb{F}_{q'}^{k'}$ such that $\sum_{i=0}^{k'-1} m_i a^i = b$. Call these set of vectors S_b . Before we proceed we recall that since $q = (q')^{k'}$, there is an isomorphism between \mathbb{F}_q and $\mathbb{F}_{q'}^{k'}$. Now note that S_b is a linear subspace and thus, for any $\gamma \in S_b$, $\prod_{c \in S_b, c \neq \gamma} (c - \gamma)$ (where we think of the operations as happening over \mathbb{F}_q) is just the product of non-zero vectors in S_b , which is some fixed constant (say) $\beta \in \mathbb{F}_q^*$. Now note that the error detection corresponding to row (a, b) is for the projected down code $\text{RS}_{S_b}[k]$. Thus, we now satisfy the second condition in Corollary 7, which implies that we can assume that the error detection can be done in linear time. Finally, for $\text{RS}_{\mathbb{F}_q}[k]$ it is well known that the second condition in Corollary 7 is also satisfied. Thus the overall runtime is bounded by $\tilde{O}(t \cdot n/\sqrt{t} + n)$, which is $\tilde{O}(ne)$, as desired.

(We remark that if we can get the best of both the random list disjunct matrix construction, i.e. $t = O(e \log n)$ and both s, s' in $\tilde{\Theta}(n/e)$, and the explicit RS code based disjunct matrix, i.e. the second condition of Corollary 7 is true, then we can have a tolerant tester with the optimal runtime of $\tilde{O}(n)$.)

The proof for the runtime for part (b) in Theorem 14 is identical and is omitted. The proof for the naive implementation runtime for part (c) is similar to part (a)—everything gets multiplied by $O(\log e)$ factor, which is at most an extra log factor. We do not know of an explicit (list) disjunct matrix that satisfies the extra requirements for part (c) and thus, we do not have any implementation with runtime better than $\tilde{O}(n^2)$.

For the rest of the section, we will focus on the other parameters of the algorithms.

All of the results above follow from a generic reduction that uses group testing. In particular, let \vec{y} be the received word that we wish to test, and \vec{c} be the nearest codeword to \vec{y} . Let $\vec{x} \in \{0, 1\}^n$ be the characteristic vector associated with error locations in \vec{y} with respect to \vec{c} . The high level idea is essentially to figure out \vec{x} using group testing.

Let M be a $t \times n$ binary matrix that is (e, e) -list disjunct. By Section 8 we can get our hands on M with $t = O(e \log n)$ with $O(e \log^2 n)$ space. Now consider the following natural algorithm.

For all $i \in [t]$, check if $\vec{y}_{M_i} \in \text{RS}_{M_i}[k]$, where M_i is the subset corresponding to the i th row of M . If so, set $r_i = 0$, else set $r_i = 1$. Run \mathcal{A} from Proposition 17 with \vec{r} as input, to get $\vec{\hat{x}}$. (STEP 1) If $\text{WT}(\vec{\hat{x}}) \geq 2e$, declare that $\geq 2e$ errors have occurred. (STEP 2) If not, declare $\leq e$ errors iff $\vec{y}_{S \setminus T} \in \text{RS}_{S \setminus T}[k]$, where T is the subset corresponding to $\vec{\hat{x}}$.

The way the algorithm is stated above, it seems to require two passes. However, using Lemma 13, we can run STEP 2 in parallel with the rest of the algorithm, resulting in a one-pass implementation.

Let \vec{z} be the result of applying M on \vec{x} . Now if it is the case that $z_i = 1$ iff $r_i = 1$, then the correctness of the algorithm above follows from the fact that M is (e, e) -list disjunct and Proposition 17. (If $\text{WT}(\vec{x}) \leq e$, then we have $S_{\vec{x}} \subseteq S_{\vec{z}}$ (where $S_{\vec{x}}$ is the subset of $[n]$ whose incidence

vector is \vec{x}) and $\text{WT}(\vec{x}) < 2e$, in which case the algorithm will declare at most e errors. Otherwise, the algorithm will “catch” the at least e errors in either STEP 1 or failing which, in STEP 2.)

However, what complicates the analysis is the fact that even though $z_i = 0$ implies $r_i = 0$, the other direction is not true. In particular, we could have $\vec{y}_{M_i} \in \text{RS}_{M_i}[k]$, even though $(\vec{y} - \vec{c})_{M_i} \neq \vec{0}$. The three parts of Theorem 14 follow from different ways of resolving this problem.

Note that if $\text{WT}(\vec{x}) \geq 2e$, then we will always catch it in STEP 2 in the worst-case. So from now on, we will assume that $0 < \text{WT}(\vec{x}) \leq e$. Let the minimum support of any row in M be s .

We begin with part (a). Let $s > k + e$ and define $\Delta = \vec{y} - \vec{c}$. Note that we are in the case where $0 < \text{WT}(\Delta) \leq e$. Since $s \geq k$ and $\vec{c}_{M_i} \in \text{RS}_{M_i}[k]$, $\vec{y}_{M_i} \in \text{RS}_{M_i}[k]$ if and only if $\Delta_{M_i} \in \text{RS}_{M_i}[k]$. Note also that for any i , $\text{WT}(\Delta_{M_i}) \leq \text{WT}(\Delta) \leq e$. Now, the distance of $\text{RS}_{M_i}[k]$ is $s - k - 1 > e$, so for every i with non-zero Δ_{M_i} , $\Delta_{M_i} \notin \text{RS}_{M_i}[k]$, which in turn means that $z_i = 1$ will always imply that $r_i = 1$ when M has the stated support. By Section 8, we have $s \geq n/(2e)$, which concludes the proof of part (a).

The following lemma follows from the random errors result in [21] and is needed for part (b):

Lemma 15 ([21]). *Let $k \leq n < q$ be integers such that $q > (\frac{n}{k})^2$. Then the following property holds for RS codes of dimension k and block length n over \mathbb{F}_q : For $\geq 1 - q^{-\Omega(k)}$ fraction of error patterns \vec{e} with $\text{WT}(\vec{e}) \leq n - 4k$ and any codeword \vec{c} , the only codeword that agrees in at least $4k$ positions with $\vec{c} + \vec{e}$ is \vec{c} .*

Now if $s \geq 4k$, then with high probability, every non-zero $\Delta_{M_i} \notin \text{RS}_{M_i}[k]$ (where Δ is as defined in the proof of part (a)). The fact that $s \geq n/(2e)$ completes the proof of part (b).

The proof of part (c) is more involved and needs a strong connection to the list decodability of the code being tested, which we discuss next.

Connection to List Decoding. Unlike the proofs of part (a) and (b) where the plain vanilla (e, e) -list disjoint matrix works, for part (c), we need and use a stronger notion of list disjoint matrices. We show that if the list disjoint matrix is picked at random, the bad tests (i.e. $r_i = 0$ even though $z_i = 1$) do not happen often and thus, one can decode the result vector even with these errors. We show that these kind of matrices suffice as long as the code being tested has good enough list decodability. The tolerant testing algorithm for a Reed-Solomon code, for instance, recursively reduces the amount of errors that need to be detected, and after application of Lemma 13, can be made to accomplish this in a single pass. We also show that the relevant list disjoint matrices can be found, with high probability, using low space and a low number of random bits.

We need to show what the forbidden subsets will be in our setting. Let C be the code we are trying to test. Let \vec{y} be the received word and let $\vec{c} \in C$ be such that $\Delta(\vec{y}, \vec{c}) \leq e \leq d/2$, where d is the distance of C . Let C be $(n - a, L + 1)$ -list decodable, that is, for any Hamming ball of radius at most $n - a$, there are at most $L + 1$ codewords from C in it. Note that if $n - a \geq \Delta(\vec{y}, \vec{c})$, then there are at most L codewords (other than \vec{c}) that agree with \vec{y} in at least a positions. Also note that each such codeword agrees with \vec{y} in at most $n - d/2$ positions. Let $T \subseteq [n]$ be the set of positions where \vec{y} and \vec{c} agree. Then define $\mathcal{F}_{a, n-d/2}(T)$ to be the (at least a) positions where codewords other than \vec{c} agree with \vec{y} . As C is $(n - a, L + 1)$ -list decodable, $|\mathcal{F}_{a, n-d/2}(T)| \leq L$.

We now show how one can use list disjoint matrices from Definition 1 to construct data stream algorithms for tolerant testing of an RS code C . Assume that there exists a $0 \leq \gamma \leq 1$, such that for every (large enough) $f \geq 1$, there exists a strongly explicit $(f, f, \gamma, \mathcal{F}_{a, n-d/2})$ -list disjoint matrix M_f . We next show how these matrices can be used to solve the tolerant testing problem for C

where we want to distinguish between the case that at most e errors have occurred and at least $2e^{\frac{2-\gamma}{1-\gamma}} + 1$ errors have occurred.

First consider the case when $e \leq k$. In this case we can use the algorithm from part (a) of Theorem 14. Thus, if $e \leq k$, then we can handle e vs. $2e$ errors, where $e \leq n/(2k)$, in space $O(e \log^2 n)$. Now consider the case when $e > k$. Let us use the matrix M_e on \vec{y} as we did before. That is, for every row of M_e (in particular, the corresponding subsets $S \subseteq [n]$), check if $\vec{y}_S \in RS_S[k]$. If so, assign $r_i = 0$ (otherwise assign $r_i = 1$). Given the result vector \vec{r} , run \mathcal{A} from Proposition 17 on it to obtain a subset $G \subseteq [n]$ such that $|G| \leq 2e$ and $[n] \setminus G$ contains at most γe errors. Thus, we have reduced the problem from at most e errors out of n positions to the problem of at most γe errors in at least $n - 2e$ positions. Then, the rest is natural: recurse on this idea. We stop when we are left with at most k errors. Note that we will need $O(\log e / \log(1/\gamma))$ many recursions. Because of these recursions, we can handle the case when there are at most e errors vs at least $2e(1 + \gamma + 2\gamma^2 + \dots) + 2k$ errors. This implies (as $e > k$) we can definitely handle e vs. $2e^{\frac{2-\gamma}{1-\gamma}} + 1$ errors. The way the idea is stated above, it seems like an $O(\log e / \log(1/\gamma))$ -pass algorithm. However, using Lemma 13, the algorithm outlined above can be implemented in one pass.

We are all done except the construction of the list disjoint matrices as defined in Definition 1:

Theorem 16. *Let $e, n, d, a, L \geq 1$ be integers such that $e \leq O\left(\frac{d}{\log L}\right)$. There exists a large enough constant $c > 1$ such that if $c \cdot e \log n \leq n$, then the following holds: There exists a $(e, e, \frac{1}{60}, \mathcal{F}_{a, n-d/2})$ -list disjoint matrix with $t = ce \log n$ rows (where for every $E \subseteq [n]$ with $|E| \leq e$, $|\mathcal{F}_{a, n-d/2}(E)| \leq L$). Further, every row has at least $\frac{n}{2e}$ ones in it. In addition, one can construct such matrices with probability at least $1 - n^{-\Omega(1)}$ using $R = O(t(e + \log L) \cdot \log e \cdot \log n)$ random bits. Further, given these R bits, any entry of the matrix can be computed in $\text{poly}(\log n)$ space.*

A folded RS code (with “folding parameter” s), is $(n - \sqrt[s+1]{sk^s n}, n^{O(s)})$ -list decodable [17]. Thus, Theorem 16 proves part (c) of Theorem 14. (Note that Theorem 16 also has the constraint that $e \leq O(d/\log L)$. However since $\log L$ is $O(\log n)$ above and as long as $d = \Omega(n)$, this bound is much weaker.)

We prove the existence of the required object by the probabilistic method. In fact this proves the second part but with $R = O(nt \log e)$. To reduce the randomness, we observe that the proof only requires bits that come from an $O(t(e + \log L) \log e)$ -wise independent source.

Let M be a $t \times n$ matrix, where each entry is chosen to be one with probability $1/e$ and $t = c \cdot e \log n$ for some large enough constant c .³

We first argue about the minimum support size of any row in M . It is easy to check that the expected Hamming weight of any row in M is exactly n/e . Thus, by the Chernoff bound, the probability that any row has Hamming weight at most $n/2e$ is upper bounded by

$$\exp\left(-\frac{n}{12e}\right) \leq \exp(-c \log n / 12) \leq n^{-190}, \quad (4)$$

where the last inequality follows for large enough c . Now by the union bound (and the fact that $t \leq n$), all the rows have Hamming weight at least $n/(2e)$ with probability at least $1 - n^{-189}$.

Next we move on to proving property (a) from Definition 1 for M with $b_1 = \frac{t}{16e}$. To this end, fix subsets $U, T \subseteq [n]$ with $|U| = e, |T| = e$ and $U \cap T = \emptyset$. Call a row $j \in [t]$ *good* if there exists a

³In this proof, we have not attempted to optimize the constants. By a conservative estimate, picking $c = 10^7$ would suffice.

$i \in U$ such that $M_{j,i} = 1$ and $M_{j,\ell} = 0$ for every $\ell \in T$. Now the probability that a row is good is exactly

$$\left(1 - \left(1 - \frac{1}{e}\right)^e\right) \left(1 - \frac{1}{e}\right)^e \geq \frac{1}{8}, \quad (5)$$

where the inequality follows if $e \geq 2$ and the fact that $(1 - 1/x)^x \leq \exp(-1) \leq 1/2$. Thus, the expected number of good rows is at least $t/8$. By the Chernoff bound, the probability that the number of good rows is at most $t/16$ is upper bounded by

$$\exp\left(-\frac{t}{96}\right) = \exp\left(-\frac{ce \log n}{96}\right) \leq n^{-190e}, \quad (6)$$

where the inequality follows for large enough c . Thus, with high probability, the number of good rows is at least $t/16$. Then by the pigeonhole principle, at least one column $i \in U$ is contained in at least $\frac{t}{16e} \stackrel{\text{def}}{=} b_2$ good rows. Taking the union bound over the $\binom{n}{e} \binom{n-e}{e}$ choices of T and U implies that with probability at least $1 - n^{-188e}$, property (a) is satisfied for every valid choice of T and U .

Next, we move on to the more involved part of the proof, which is to prove property (b) in Definition 1. To this end, given any $T \subseteq [n]$ with $|T| \leq e$ and a column $i \in [n]$ we will upper bound the probability that at least b_1 tests that contain i are themselves contained in some subset in $\mathcal{F}_{a,n-d/2}(T)$. It turns out that this probability will be $n^{-O(1)}$, which is not enough to apply the union bound over all the $\binom{n}{e}$ choices of T . We then observe that these probabilities are almost independent for any $\Omega(e)$ such columns, which is sufficient for the union bound over all choice of T to go through.

Fix a subset $T \subseteq [n]$ with $|T| \leq e$ and a subset $S \in \mathcal{F}_{a,n-d/2}(T)$. (Note that $|S| \leq n - d/2$.) We say that a row $j \in [t]$ *avoids* S ($\mathcal{F}_{a,n-d/2}(T)$ resp.) if the j th row (which we will denote by $M(j)$) is not a subset of S (any subset in $\mathcal{F}_{a,n-d/2}(T)$ resp.). In other words, if j avoids S then $M_{j,i} = 1$ for some $i \notin S$. Thus, we have

$$\Pr[j \text{ doesn't avoid } S] = \left(1 - \frac{1}{e}\right)^{n-|S|} \leq \left(1 - \frac{1}{e}\right)^{-d/2} \leq \exp\left(-\frac{d}{2e}\right), \quad (7)$$

where first inequality follows from the fact that $|S| \leq n - d/2$.

Fix a column $i \in [t]$. Note that if $i \notin S$ and $M(j)$ contains i , then j does avoid S . Now, if $i \in S$ and given that the probability in (7) only depends on the indices $i \notin S$, we get that $\Pr[j \text{ contains } i \text{ and doesn't avoid } S] \leq \exp(-d/(2e))/e$. Thus, whether $i \in S$ or not, we have by the union bound

$$\Pr[j \text{ contains } i \text{ and doesn't avoid } \mathcal{F}_{a,n-d/2}(T)] \leq \frac{L}{e} \cdot \exp\left(-\frac{d}{2e}\right) \leq \frac{1}{80e}, \quad (8)$$

where the last inequality follows if $e \leq \frac{d}{10 \ln L}$.

Now call a row j *i-bad* if it contains i but does not avoid $\mathcal{F}_{a,n-d/2}(T)$. (If it contains i and avoids $\mathcal{F}_{a,n-d/2}(T)$, then call it *i-good*.) Note that we need to show that for at least $(1 - \gamma)d$ columns $i \in T$, there are at most b_1 *i-bad* rows. Thus, by (8), the expected number of *i-bad* rows is at most $t/(80e)$, or the expected number of *i-good* rows is at least $79t/(80e)$. By the Chernoff bound, we have

$$\Pr\left[\text{Number of } i\text{-good rows} < \frac{78t}{80e} = \frac{39t}{40e}\right] \leq \exp\left(-\frac{t}{3 \cdot 79 \cdot 80e}\right) \leq n^{-190}, \quad (9)$$

where the last inequality follows from large enough c . Since the expected Hamming weight of any column is t/e , the Chernoff bound implies that

$$\Pr \left[\text{Column } i \text{ has Hamming weight } \geq \frac{81t}{80e} \right] \leq \exp \left(-\frac{t}{3 \cdot 80^2 e} \right) \leq n^{-190}, \quad (10)$$

where again the last inequality follows for large enough c . Thus, (9) and (10) imply that

$$\Pr \left[\text{Number of } i\text{-bad rows} > \frac{3t}{80e} \right] \leq 2 \cdot n^{-190} \leq n^{-189}, \quad (11)$$

where the last inequality is true for $n \geq 2$. Unfortunately, the bound above is too weak to apply the union bound over all the $\binom{n}{e}$ choices of T . However, we get around this obstacle by proving that for any $\Omega(e)$ values of $i \in [t]$, the probabilities above are essentially independent.

Call a column $i \in T$ *bad* if the number of bad i -rows is at least b_1 (for some $\frac{3t}{80e} < b_1 < \frac{t}{16e}$ to be fixed later). For notational convenience, define $\ell = \frac{e}{60}$. Next we are going to show that for any subset $V = \{i_1, \dots, i_\ell\} \subseteq T$,

$$\Pr [\text{Every } j \in V \text{ is bad}] \leq n^{-3e}. \quad (12)$$

Note that the above implies that the probability that more than ℓ columns in T are bad is upper bounded by

$$\binom{e}{\frac{e}{60}} n^{-3e} \leq n^{-2e},$$

where the last inequality follows for $n \geq 180$. Thus, the probability that for some $T \subseteq [n]$ with $|T| \leq e$, there are more than ℓ bad columns in T , by the union bound, is upper bounded by n^{-e} . Thus, property (b) is true with probability at least $1 - n^{-e}$.

To complete the proof, we will prove (12). Note that we can rewrite the probability in (12) as

$$\Pr \left[i_\ell \text{ is bad} \mid \bigwedge_{j \in V \setminus i_\ell} j \text{ is bad} \right] \cdot \prod_{j \in V \setminus i_\ell} \Pr [j \text{ is bad}] \leq \Pr \left[i_\ell \text{ is bad} \mid \bigwedge_{j \in V \setminus i_\ell} j \text{ is bad} \right] \cdot n^{-2e},$$

where the inequality follows from (11) and the fact that $b_1 > 3t/(80e)$. Thus, we will be done if we can show

$$\Pr \left[i_\ell \text{ is bad} \mid \bigwedge_{j \in V \setminus i_\ell} j \text{ is bad} \right] \leq n^{-e}. \quad (13)$$

To this end, let $B \subseteq [t]$ be the rows that contain at least one column from $V \setminus i_\ell$. Note that

$$\mathbf{E}[|B|] = t \left(1 - \left(1 - \frac{1}{e} \right)^\ell \right). \quad (14)$$

By the Chernoff bound, we have

$$\Pr \left[|B| \geq \frac{6t}{5} \left(1 - \left(1 - \frac{1}{e} \right)^\ell \right) \right] \leq \exp \left(-\frac{t}{75} \cdot \left(1 - \left(1 - \frac{1}{e} \right)^\ell \right) \right). \quad (15)$$

As for any real $x > 0$ and integer $y > 0$ with $xy < 1$, $1 - xy \leq (1 - x)^y \leq 1 - xy + (xy)^2/2$, we have $59/60 \leq (1 - \frac{1}{e})^\ell \leq 59/60 + 1/7200 < 119/120$. This along with (15) implies that

$$\Pr \left[|B| > \frac{t}{50} \right] \leq \exp \left(-\frac{t}{75 \cdot 120} \right) \leq n^{-190e},$$

where the last inequality follows for large enough c . Similarly, one can show that

$$\Pr \left[\text{Number of rows in } B \text{ that contain } i_\ell > \frac{t}{50e} \right] \leq n^{-190}.$$

We do a conservative estimate and assume that all tests in B that contain i_ℓ are i_ℓ -bad. Because of the bound above, w.l.o.g. with all but an n^{-190} probability, we can assume that $|B| = t/50$. Using the same calculation⁴ as we did to obtain (11), we can show that

$$\Pr \left[\text{Number of } i_\ell\text{-bad rows in } B > \frac{3 \cdot 49t}{4000e} \right] \leq 2 \cdot n^{-49 \cdot 190/50} \leq n^{-185}.$$

Adding in the number of rows in B that contain i_ℓ , we obtain that

$$\Pr \left[\text{Number of } i_\ell\text{-bad rows} > \frac{t}{50e} + \frac{147t}{4000e} \mid \bigwedge_{j \in V \setminus i_\ell} j \text{ is bad} \right] \leq n^{-190} + n^{-185} \leq n^{-180}.$$

Picking $b_1 = \frac{t}{e} \left(\frac{1}{50} + \frac{147}{4000} \right) < b_2$ completes the proof of (13). Thus, we have completed the proof of the existence of the desired $(e, e, \frac{1}{60}, \mathcal{F}_{a, n-d/2})$ -list disjoint matrix.

In fact, the proof shows that the required matrices can be computed with high probability. However, at least $\Omega(n)$ random bits are required, which is too high for any data stream application. Next we point out that the proof only requires limited independence and hence, we can get away with much fewer random bits. In the remainder of the proof, we will think of the bits of M to come from some k -wise independent source that contain bit strings of length tn .

We now go through the proof above and estimate the amount of independence needed. The first place that needs independence is (4) and we claim that $O(\log n)$ -wise independence suffices. This follows from the tail bounds for k -wise independent sources from [6]. In particular, Bellare and Rompel show that for a k -wise independent source, the sum of binary random variables with mean μ can have a deviation of strictly more than A with probability at most $8 \cdot \left(\frac{\mu k + k^2}{A^2} \right)^{k/2}$. Note that in our case $\mu = n/e$, $A = n/(2e)$ and since $n/e \geq c \log n$, picking a $O(\log n)$ -wise independent source works.

The next places in the proof that use independence are (5) and (6). It is easy to check that the calculations go through if we have $2et$ -wise independence. Next, independence is used in (7). Note that in this case we need to upper bound the probability by $(1 - 1/e)^{\Omega(e \log L)}$. Thus, picking $O(e \log L)$ -wise independence works for this case. Next (9) needs t/e -wise independence. This follows from the tail bound for k -wise independence from [6]. Note that we actually need the product of the independence used in (7) and (9), that is, we need a total of $O(t \log L)$ -wise independence. For (10) t -wise independence suffices. Finally for (14) and (15) we need ℓt -wise independence. In fact, again using the bound from [6], we can get away with $O(t)$ -wise independence.

⁴We need to replace t by $49t/50$. Further in (7), we need to replace $n - |S|$ by $n - |S| - e/60$ as in the worst case $\{i_1, \dots, i_{\ell-1}\} \subseteq [n] \setminus S$. However, this does not change the upper bound in (8) as long as we pick $e \leq d/(15 \ln L)$.

Thus, overall we need $O(t(e + \log L))$ -wise independence. Generally, k -wise independent sources are for unbiased bits where as we need random bits that take a value of one with probability $1/e$. However, since we can get such random bits from $O(\log e)$ unbiased bits, we will need $O(t(e + \log L) \log e)$ -wise independent sources containing $O(nt \log e)$ bit strings. Using well-known construction of k -wise independent sources, we can get away with $R = O(t(e + \log L) \log e \cdot \log(nt))$ random bits. This completes the proof as $t \leq n$.

6.1 List Disjunct Matrices.

We begin with the definition of a stronger kind of list disjunct matrices:

Definition 1. Let $n, s_1, s_2, e, \ell, L \geq 1$ be integers with $s_1 \leq s_2$ and let $0 \leq \gamma \leq 1$ be a real. For any subset $T \subset [n]$ such that $|T| \leq e$, let $\mathcal{F}_{s_1, s_2}(T)$ be a collection of forbidden subsets of $[n]$ of size in the range $[s_1, s_2]$ such that $|\mathcal{F}_{s_1, s_2}(T)| \leq L$. A $t \times n$ binary matrix M is called a $(e, \ell, \gamma, \mathcal{F}_{s_1, s_2})$ -list disjunct matrix if there exist integers $0 \leq b_1 < b_2$ such that the following hold for any $T \subseteq [n]$ with $|T| \leq e$:

1. For any subset $U \subseteq [n]$ such that $|U| \geq \ell$ and $U \cap T = \emptyset$, there exists an $i \in U$ with the following property: The number of rows where the i th column of M has a one and all the columns in T have a zero is at least b_2 .
2. The following holds for at least $(1 - \gamma)e$ many $i \in T$: Let R_i denote all the rows of M (thought of as subsets of $[n]$) that contain i . Then $|\{U \in R_i \mid U \subseteq V, \text{ for some } V \in \mathcal{F}_{s_1, s_2}(T)\}| \leq b_1$.

The definition might appear complicated but it is setup to easily imply Proposition 17. Further, a $(e, \ell, 0, \emptyset)$ -list disjunct matrix (with $b_1 = 0$ and $b_2 = 1$) is the same as the (e, ℓ) -list disjunct matrix considered in [18]. Further, an $(e, 1)$ -list disjunct matrix is the well-known e -disjunct matrix [13].

Let us also define the following error version of group testing that will be relevant to our scenario.

Definition 2. Let $n, s_1, s_2, e, L \geq 1$ be integers. For every $T \subseteq [n]$ such that $|T| \leq e$, let $\mathcal{F}_{s_1, s_2}(T)$ be the collection of forbidden subsets as in Definition 1. Then $(e, \mathcal{F}_{s_1, s_2})$ -group testing works in the following manner: Given a set of defectives $T \subseteq [n]$ such that $|T| \leq e$, any test $U \subseteq [n]$ behaves as follows: If $U \cap T = \emptyset$, then the test will return an answer of 0. If $U \cap T \neq \emptyset$ and $U \subseteq V$ for some $V \in \mathcal{F}_{s_1, s_2}(T)$, then the test will return an answer of 0. Otherwise the test returns an answer of 1.

The algorithm \mathcal{A} in the below proposition is a natural generalization of the standard decoding algorithm for e -disjunct matrices [13].

Proposition 17. Let $n, e, \ell, s_1, s_2, L, \gamma, \mathcal{F}_{s_1, s_2}$ be as in Definition 1. Let M be a $(e, \ell, \gamma, \mathcal{F}_{s_1, s_2})$ -list disjunct matrix with t rows. Finally, consider an outcome vector \vec{r} of applying M to a set of defectives E with $|E| \leq e$ in the $(e, \mathcal{F}_{s_1, s_2})$ -group testing scenario. Then there exists an algorithm \mathcal{A} , which given \vec{r} can compute a set G such that $|G| \leq \ell + e - 1$ and $|E \setminus G| \leq \gamma e$. Further, \mathcal{A} uses $O(t + \log n + S(t, n))$ space, where $S(t, n)$ is the space required to compute any entry of M .

6.2 Proof of Proposition 17

The algorithm \mathcal{A} is very simple: Go through every column $i \in [n]$ and declare $i \notin G$ if and only if the number of rows $j \in [t]$ where $M_{j,i} = 1$ but $r_j = 0$ is at least b_2 . It is easy to check that \mathcal{A} has the claimed space requirement. The correctness of \mathcal{A} follows from Definitions 1 and 2. To see this

note that if $|G| \geq e + \ell$, i.e. $|G \setminus E| \geq \ell$, then by part (a) of Definition 1, there exists an $i \in G \setminus E$ with the following property: There are at least b_2 rows $j \in [t]$ such that $M_{j,i} = 1$ but $M_{j,i'} = 0$ for every $i' \in E$. By Definition 2, for every such j , $r_j = 0$. Thus, by definition of \mathcal{A} , i cannot be in G . Now consider an $i \in E$ for which property (b) of Definition 1 holds. Now by Definition 2, there are at most b_1 rows $j \in [t]$ such that $M_{j,i} = 1$ and $r_j = 0$. Since $b_1 < b_2$, \mathcal{A} includes i in G . This implies that $|E \setminus G| \leq \gamma e$.

The space requirement of $O(e^2 \log^2 n)$ of part (c) is unsatisfactory. Reducing the amount of randomness needed to something like $O(e \log n)$ will realize the full potential of our algorithm. We leave this as an open problem.

7 Limitations of our techniques

One shortcoming of Theorem 14 is that to distinguish between (say) at most e and at least $2e$ errors, we needed $e \cdot s \leq O(n)$, where s is the minimum support size of any test. Another shortcoming is that we need $O(e \log n)$ space. In this section, we prove that our techniques cannot overcome these limits.

We begin with some quick notation. For any $k \geq 1$, a k^{++} query to a string $x \in \{0, 1\}^n$ corresponds to a subset $S \subseteq [n]$. The answer to the query is x_S if $\text{WT}(x_S) < k$, otherwise the answer is k^{++} (signifying that $\text{WT}(x_S) \geq k$). (This is a natural generalization of k^+ decision trees considered by Aspnes et al. [5].) A k^{++} algorithm to solve the (ℓ, t, n) -threshold function makes a sequence of k^{++} queries to the input $x \in \{0, 1\}^n$, and can tell whether $\text{WT}(x) \leq \ell$ or $\text{WT}(x) \geq t$. If we think of x as being the indicator vector for error locations, then our reduction from tolerant testing to error detection can be thought of as a 1^{++} algorithm for the $(e, O(e))$ -threshold function.

First we show that the minimum support size that we obtain in our reduction, even with the stronger k^{++} primitive, is nearly optimal.

Theorem 18. *Let $0 \leq \ell < t \leq n$ and $k \geq 1$ be integers. Let $\varepsilon < 1/2$ be a constant real. Then any non-adaptive, randomized k^{++} algorithm for the (ℓ, t, n) -threshold problem with error probability at most ε , where all the queries have support size at least s , needs to make at least $e^{s\ell/n}/n^{O(k)}$ queries. In particular, any algorithm that makes a sublinear number of queries needs to satisfy $s \cdot \ell \leq O(kn \log n)$.*

7.1 Proof of Theorem 18

Define the following distribution \mathcal{D} on inputs in $\{0, 1\}^n$: uniformly distribute a probability mass of $1/2$ over the $\binom{n}{\ell}$ vectors of Hamming weight exactly ℓ (call this set \mathcal{N}) and the rest of the probability mass uniformly over the $\binom{n}{t}$ vectors of Hamming weight t (call this set \mathcal{Y}). We will show that any deterministic non-adaptive k^{++} algorithm with an error probability at most ε (according to \mathcal{D}) must make at least $\frac{e^{s\ell/n}}{n^{O(k)}}$ queries. Yao's lemma will then complete the proof.

Fix an arbitrary k^{++} algorithm A that has error probability at most ε . Thus, A outputs the correct value of 0 in at least $\frac{1/2}{(1/2-\varepsilon)} \geq 1 - 2\varepsilon$ fraction of elements in \mathcal{N} (call this set of elements \mathcal{N}_0). Similarly, the algorithm outputs the correct value of 1 in at least $1 - 2\varepsilon$ fraction of the elements in \mathcal{Y} (call this set \mathcal{Y}_1). Any k^{++} query is said to *cover* a pair of inputs $(x_0, x_1) \in \mathcal{N}_0 \times \mathcal{Y}_1$, if it outputs different answers for the inputs x_0 and x_1 . Note that all the pairs in $\mathcal{N}_0 \times \mathcal{Y}_1$ have to be covered by some query in A .

To complete the proof, we will show that at least $e^{s\ell/n}/n^{O(k)}$ queries are needed to cover $\mathcal{N}_0 \times \mathcal{Y}_1$. To this end given an arbitrary query Q of support $i \geq s$, we will bound the number of pairs it can cover (call this number of pairs P_Q). Note that Q will not cover a pair (x_0, x_1) if both x_0 and x_1 have at least k ones in the support of Q . Thus, to upper bound P_Q , we will count the number of pairs (x_0, x_1) such that either x_0 or x_1 have support $< k$ in Q . This latter count is clearly upper bounded by

$$\max \left(\binom{n}{\ell} \cdot \left(\sum_{j=0}^{k-1} \binom{i}{j} \binom{t}{j} \binom{n-i}{t-j} \right), \left(\sum_{j=0}^{k-1} \binom{i}{j} \binom{\ell}{j} \binom{n-i}{\ell-j} \right) \binom{n}{t} \right),$$

where for notational convenience we define $\binom{a}{b} = 0$ for $b > a$. We claim that the above is upper bounded by (see Appendix A for a proof):

$$kn^{3(k-1)} \cdot \max \left(\binom{n}{\ell} \binom{n-s}{t'}, \binom{n-s}{\ell'} \binom{n}{t} \right),$$

where $t' = \max_{0 \leq j \leq k-1} \{t-j | t-j \leq n-s\}$ and $\ell' = \max_{0 \leq j \leq k-1} \{\ell-j | \ell-j \leq n-s\}$. The way we are going to proceed with the rest of the proof, the maximum in the above will occur for the second argument, i.e. from now on, we have that for any query Q ,

$$P_Q \leq kn^{3(k-1)} \binom{n-s}{\ell'} \binom{n}{t} \stackrel{\text{def}}{=} P_{\max}. \quad (16)$$

As

$$|\mathcal{N}_0 \times \mathcal{Y}_1| \geq (1-2\varepsilon)^2 \binom{n}{\ell} \binom{n}{t}, \quad (17)$$

by the pigeonhole principle, the number of queries that A needs to make is at least

$$\frac{|\mathcal{N}_0 \times \mathcal{Y}_1|}{P_{\max}} \geq \frac{(1-2\varepsilon)^2 \binom{n}{\ell} \binom{n}{t}}{kn^{3(k-1)} \binom{n-s}{\ell'} \binom{n}{t}} \quad (18)$$

$$\geq \frac{(1-2\varepsilon)^2 \binom{n}{\ell'}}{kn^{4(k-1)} \binom{n-s}{\ell'}} \quad (19)$$

$$\geq \frac{\sqrt{8}(1-2\varepsilon)^2 e^{s\ell'/n}}{kn^{4(k-1)} \sqrt{27(n+1)}} \quad (20)$$

$$\geq \frac{\sqrt{8}(1-2\varepsilon)^2 e^{s\ell/n}}{kn^{4(k-1)} e^{k-1} \sqrt{27(n+1)}}. \quad (21)$$

In the above, (18) follows from (17) and (16). (19) follows from the following argument. Note that if $\ell < n/2$ then $\binom{n}{\ell} \geq \binom{n}{\ell'}$. If $\ell' > n/2$, then $\binom{n}{\ell} \geq \binom{n}{\ell}/n^{\ell-\ell'}$. Finally if $\ell' < n/2$ and $\ell \geq n/2$, then $\binom{n}{\ell} \geq \binom{n}{\ell'}$ if $|n/2 - \ell| < |n/2 - \ell'|$ otherwise $\binom{n}{\ell} \geq \binom{n}{\ell'}/n^{|n/2-\ell|-|n/2-\ell'|}$. Thus, in all cases, $\binom{n}{\ell} \geq \binom{n}{\ell'}/n^{\ell-\ell'} \geq \binom{n}{\ell'}/n^{k-1}$, where the last inequality follows from the fact that $\ell - \ell' \leq k-1$. (20) follows from Lemma 19. Finally (21) follows from the fact that $\ell' \geq \ell - k + 1$ and $s \leq n$.

We are done except for the following lemma:

Lemma 19. *Let $a \leq n$ and $b \leq n-a$ be integers. Then*

$$\frac{\binom{n}{b}}{\binom{n-a}{b}} \geq e^{ab/n} \cdot \sqrt{\frac{8}{27(n+1)}}.$$

Proof. Stirling's approximation can be used to obtain the following bound for $n \geq 1$,

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n \leq n! \leq \sqrt{3\pi n} \left(\frac{n}{e}\right)^n.$$

In particular, this implies that for any $x \leq y$,

$$\frac{1}{3} \sqrt{\frac{2y}{\pi x(y-x)}} \cdot \frac{y^y}{x^x(y-x)^{y-x}} \leq \binom{y}{x} \leq \frac{1}{2} \sqrt{\frac{3y}{\pi x(y-x)}} \cdot \frac{y^y}{x^x(y-x)^{y-x}}.$$

Using the bound above, we get

$$\frac{\binom{n}{b}}{\binom{n-a}{b}} \geq f(a, b, n) \cdot \frac{n^n(n-a-b)^{n-a-b}}{(n-b)^{n-b}(n-a)^{n-a}}, \quad (22)$$

where

$$f(a, b, n) = \sqrt{\frac{8}{27} \cdot \frac{n(n-a-b)}{(n-b)(n-a)}} \geq \sqrt{\frac{8}{27 \left(1 + \frac{ab}{n(n-a-b)}\right)}} \geq \sqrt{\frac{8}{27(1+n)}}, \quad (23)$$

where the last inequality used the facts that $ab \leq n^2$ and $n-a-b \geq 1$.

Now consider the following sequence of relationships

$$\begin{aligned} \frac{n^n(n-a-b)^{n-a-b}}{(n-b)^{n-b}(n-a)^{n-a}} &= \left(\frac{n}{n-a}\right)^b \left(\frac{n}{n-b}\right)^a \left(\frac{n(n-a-b)}{(n-a)(n-b)}\right)^{n-a-b} \\ &= \frac{1}{\left(1 - \frac{a}{n}\right)^b} \cdot \frac{1}{\left(1 - \frac{b}{n}\right)^a} \cdot \frac{1}{\left(1 + \frac{ab}{n(n-a-b)}\right)^{n-a-b}} \\ &\geq \frac{1}{e^{-ab/n}} \cdot \frac{1}{e^{-ab/n}} \cdot \frac{1}{e^{ab/n}} \\ &= e^{ab/n}, \end{aligned} \quad (24)$$

where the inequality follows from the following two facts (for $x, y > 0$):

$$\left(1 + \frac{x}{y}\right)^y \leq e^x \text{ and } (1-x)^y \leq e^{-xy}.$$

(22), (23) and (24) complete the proof. ■

Note that our reduction maps one tolerant testing problem instance (where say we want to distinguish between at most e error vs. at least $2e$ errors) to $O(e \log n)$ many instances of error detection. Next we show that this is essentially unavoidable even if we use k^{++} queries for constant k . The following result follows from the results in [5]:

Theorem 20. *Let $0 \leq \ell < t \leq n$ and $k \geq 1$ be integers. Then any adaptive, deterministic k^{++} algorithm for the (ℓ, t, n) -threshold problem makes $\Omega(\ell/k)$ queries.*

7.2 Proof of Theorem 20

The proof will be by an adversarial argument to show that if $r < \ell/k$ k^{++} queries are made then there exist two inputs \vec{x} and \vec{y} on which the answers to the queries will be the same, yet $\text{WT}(\vec{x}) \leq \ell$ and $\text{WT}(\vec{y}) \geq t$. Note that the existence of such a pair of inputs will complete our proof.

We will think of the adversary as maintaining a set of positions $U(i)$ after the i th step. The invariance that the adversary will maintain is that $U(i-1) \subseteq U(i)$ and more importantly, that any input \vec{x} such that $\vec{x}_{U(r)} = \vec{1}$ will be consistent with answers to the queries. Finally, it is also the case that $|U(i)| \leq ki$. Note that if we can come up with a way to construct these subsets $\{U(i)\}_{i=1}^r$, then the proof will be done (consider the inputs $\vec{1}$ and \vec{a} such that $\vec{a}_{U(r)} = \vec{1}$ and $\vec{a}_{[n] \setminus U(r)} = \vec{0}$).

To complete the proof, we will show how the adversary can construct the set $U(i)$. Given the i th k^{++} query $S \subseteq [n]$, the adversary constructs $U(i)$ as follows: Let $S' = S \setminus U(i-1)$. If $|S'| \leq k$, then let $U(i) = U(i-1) \cup S$. Otherwise pick an arbitrary subset $T \subseteq S'$ such that $|T| = k$ and define $U(i) = U(i-1) \cup T$. In both cases, the adversary answers the query as follows: If $|S| \geq k$, return an answer of k^{++} , otherwise report that the substring indexed by S is the all ones vector. It is easy to check that $U(i)$ satisfies all the required properties.

8 Randomness Efficient Construction of List Disjunct Matrices

In this section, we show that (e, e) -list disjunct matrices can be constructed with $t = O(e \log n)$ rows (each with support at least $n/(2e)$) with $O(e \log^2 n)$ random bits.

For this, we will need Nisan's PRG for space bounded computation [20]. Nisan's result states that there exists a function $G : \{0, 1\}^T \rightarrow \{0, 1\}^R$ such that any Finite State Machine that uses $O(S)$ space and R random bits, cannot distinguish between truly unbiased random R bits and the bits $G(x)$ (for x chosen randomly from $\{0, 1\}^T$) for $T = O(S \log R)$ with probability more than $2^{-O(S)}$. Further, any bit of $G(x)$ can be computed, given the T random bits x (and $O(S)$ extra storage).

We first use the probabilistic method to show that the required object exists with high probability. Then we show that the proof can be implemented in low space and use Nisan's PRG to complete the proof.

Let $t = c \cdot e \log n$, where c is some large enough constant so that all calculations go through. Also let $\alpha \geq 1$ be a large enough constant to be determined later. We will also assume that $t \leq n$ so that $n/e \geq c \log n$. Let M be a random $t \times n$ matrix, where each entry is one independently with probability $1/e$. Now to prove that M has the required property, we show that it satisfies the following two properties with high probability:

- (a) Every row of M has Hamming weight at least $\frac{n}{2e}$.
- (b) For any two disjoint subsets $S, T \subseteq [n]$ such that $|S| = |T| = e$, there is at least one row such that at least one column in T has a one in it while all the columns in S have a zero in it.

We begin with (a). Note that in expectation any row has n/e ones in it. Thus, by Chernoff bound, the probability that any row has Hamming weight at most $n/(2e)$ is upper bounded by

$$\exp\left(-\frac{n}{12e}\right) \leq \exp\left(-\frac{c}{12} \cdot \log n\right) \leq n^{-2\alpha},$$

where the last inequality follows for $c \geq 24\alpha$ and the first inequality follows from the assumption that $t \leq n$.

Next, we move to (b). Fix a row $j \in [t]$. Now the probability that $\forall_{i \in T} M_{j,i} = 1$ and $\forall_{i \in S} M_{j,i} = 0$ is exactly

$$\left(1 - \left(1 - \frac{1}{e}\right)^e\right) \left(1 - \frac{1}{e}\right)^e \geq \frac{1}{8},$$

where the last inequality follows for $e \geq 2$. Thus, the probability that there does not exist a row as desired in part (b) is upper bounded by

$$\left(\frac{7}{8}\right)^{ce \log n} \leq n^{-(2+\alpha)e},$$

where the last inequality follows for $c \geq 20(2 + \alpha)$. Thus, by the union bound, part (a) does not hold with probability at most $n^{-2\alpha}$ (for $e \geq 2$).

Thus, M does not have the desired property with probability at most $n^{-\alpha}$ (for $n \geq 2$).

Next, we estimate the space required to implement the proof above, i.e. given $R = nt$ bits of the entries in M , we need to figure out how much space is needed to verify whether M has the required property or not. For part (a), we need $O(\log t + \log n)$ space to keep track of the row and $O(\log(n/e))$ bits to check if the row has Hamming weight at least $n/(2e)$. So we can implement part (a) with $O(\log n)$ space. For part (b), we need $O(e \log n)$ space to keep track of the subsets S and T . For given S and T , we need $O(\log t)$ space to keep track of the rows and $O(\log n)$ space to verify if it is the row that “takes care” of S and T . Thus, for part (b) we need $O(e \log n)$ space.

Thus, overall we have $S = O(e \log n)$. We are almost done, except for one small catch: Nisan’s PRG deals with unbiased bits but we need random bits that are biased. However, we can obtain a random bit that is one with probability $1/e$ from $O(\log e)$ unbiased bits (by declaring the final bit to be one if and only if all the unbiased bits are 1). Thus, we can convert the proof above to use $R' = O(\log e \cdot R)$ unbiased random bits. Further, this conversion needs an extra $O(\log \log e + \log R)$ space, which implies that the total space used is $S' = O(e \log n)$.

Thus, by Nisan’s PRG we would be done with $O(S' \log R') = O(e \log^2 n)$ random bits. Using Nisan’s PRG will increase the error probability to $n^{-\alpha} + 2^{-O(S)}$, which can be made to be polynomially small by picking α appropriately.

Acknowledgments

We thank Venkat Guruswami, Steve Li and Ram Swaminathan for helpful discussions. Thanks to Chris Umans for pointing out [6] to us.

References

- [1] L. M. Adleman and H. W. Lenstra. Finding irreducible polynomials over finite fields. In *STOC '86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 350–355, New York, NY, USA, 1986. ACM.
- [2] M. Alekhnovich. Linear diophantine equations over polynomials and soft decoding of reed-solomon codes. *IEEE Transactions on Information Theory*, 51(7):2257–2265, 2005.
- [3] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.

- [4] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [5] J. Aspnes, E. Blais, M. Demirbas, R. O’Donnell, A. Rudra, and S. Uurtamo. k+ decision trees, 2010. Manuscript.
- [6] M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 276–287, 1994.
- [7] E. Ben-Sasson, V. Guruswami, T. Kaufman, M. Sudan, and M. Viderman. Locally testable codes require redundant testers. In *IEEE Conference on Computational Complexity*, pages 52–61, 2009.
- [8] E. Ben-Sasson, P. Harsha, and S. Raskhodnikova. Some 3cnf properties are hard to test. *SIAM J. Comput.*, 35(1):1–21, 2005.
- [9] I. F. Blake, S. Gao, A. J. M. (Editor), R. C. Mulin, S. A. Vanstone, and T. Yaghoobian. *Applications of Finite Fields*. Kluwer Academic Publishers, 1993.
- [10] C. L. Chen and M. Y. Hsiao. Error-correcting codes for semiconductor memory applications: A state-of-the-art review. *IBM Journal of Research and Development*, 28(2):124–134, 1984.
- [11] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson. RAID: High-performance, reliable secondary storage. *ACM Computing Surveys*, 26(2):145–185, 1994.
- [12] I. Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.
- [13] D.-Z. Du and F. K. Hwang. *Combinatorial Group Testing and its Applications*. World Scientific, 2000.
- [14] J. Elerath. Hard-disk drives: The good, the bad, and the ugly. *Communications of the ACM*, 52(6):38–45, 2009.
- [15] Z. Füredi. On r -cover-free families. *J. Comb. Theory, Ser. A*, 73(1):172–173, 1996.
- [16] V. Guruswami and A. Rudra. Tolerant locally testable codes. In *Proceedings of the 9th International Workshop on Randomization and Computation (RANDOM)*, pages 306–317, 2005.
- [17] V. Guruswami and A. Rudra. Explicit codes achieving list decoding capacity: Error-correction up to the Singleton bound. *IEEE Transactions on Information Theory*, 54(1):135–150, January 2008.
- [18] P. Indyk, H. Q. Ngo, and A. Rudra. Efficiently decodable non-adaptive group testing. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1126–1142, 2010.
- [19] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [20] N. Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

- [21] A. Rudra and S. Uurtamo. Two theorems in list decoding. *ECCC Technical Report TR10-007*, 2010.
- [22] V. Shoup. *A Computational Introduction to Number Theory and Algebra*. Cambridge University Press, 2008. 2nd Edition.
- [23] D. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6):1723–1732, 1996.
- [24] M. Sudan. Algorithmic introduction to coding theory, 2001. Lecture Notes available at <http://people.csail.mit.edu/madhu/FT01/>.

A Upper bounding a sum

We begin with the sum

$$\sum_{j=0}^{k-1} \binom{i}{j} \binom{t}{j} \binom{n-i}{t-j}.$$

Let s, i, t and n be such that $s \leq i \leq n$ and $t \leq n$. The sum above is then upper bounded by

$$\sum_{j=0}^{k-1} \binom{n}{j} \binom{n}{j} \binom{n-s}{t-j}.$$

Now define j^* to be the minimum $0 \leq j \leq k-1$ such that $t - j^* \leq n - s$ (if no such j^* exists then the sum is 0). Now upper bounding $\binom{n}{j} \leq n^{k-1}$ for $j \leq k-1$, we can again upper bound the sum above by

$$n^{2(k-1)} \sum_{j=j^*}^{k-1} \binom{n-s}{t-j}.$$

From the bound that $\binom{a}{b} \leq a \binom{a}{b-1}$, we get that $\binom{n-s}{t-j} \leq (n-s)^{j-j^*} \binom{n-s}{t-j^*} \leq n^{k-1} \binom{n-s}{t-j^*}$. This along with the bound above implies that

$$\sum_{j=0}^{k-1} \binom{i}{j} \binom{t}{j} \binom{n-i}{t-j} \leq kn^{3(k-1)} \binom{n-s}{t'},$$

where $t' = t - j^*$, as desired. Similarly one can show that

$$\sum_{j=0}^{k-1} \binom{i}{j} \binom{\ell}{j} \binom{n-i}{\ell-j} \leq kn^{3(k-1)} \binom{n-s}{\ell'},$$

where $\ell' = \max_{0 \leq j \leq k-1} \{\ell - j | \ell - j \leq n - s\}$.